

**Contributions to Data Quality in IAM:
Assessment, Improvement and Application**

Dissertation zur Erlangung des Grades eines
Doktors der Wirtschaftswissenschaft



eingereicht an der
Fakultät für Wirtschaftswissenschaften
der Universität Regensburg

vorgelegt von:

Sascha Kern

Berichterstatter:

Prof. Dr. Günther Pernul

Prof. Dr. Bernd Heinrich

Tag der Disputation: 18.12.2025

Abstract

Cybersecurity breaches pose a serious threat to society, potentially causing data leaks, financial losses, and disruptions to critical infrastructure. Identity and Access Management (IAM), a pillar of IT security management, aims to minimize potential attack surfaces. Despite its importance, many organizations struggle to implement effective IAM. An important success factor is data quality: While sufficient data quality is a prerequisite for enabling effective access control, low data quality leads to errors, causing security vulnerabilities and operational costs. This dissertation addresses this problem with research on the assessment and improvement of data quality in IAM. The scope of this research is the quality of access control policies and attributes, with focus areas including quality assessment, quality maintenance, and access reviews. The results were made available in seven peer-reviewed publications, which are part of this cumulative dissertation.

Acknowledgement

Five years is a long time. I received support from a lot of people while working on this dissertation, be it directly or indirectly. First of all, I want to thank my supervisor Professor Dr. Günther Pernul, who accompanied me throughout this journey. Whether it was about the general direction of this dissertation, specific work results, or the preparation of a publication just before the submission deadline: I could always rely on your well-meaning but honest counsel. I would also like to thank my second supervisor Professor Dr. Bernd Heinrich for his helpful advice. I would like to thank my co-authors: the tireless Thomas Baumer, with whom I spent many days and nights conducting research. Our joint conference trips were the greatest reward for this work. Dr. Sebastian Groll, who invited me to join his research right at the beginning of my dissertation. Our fruitful discussions and your enthusiasm always inspired me. Dr. Ludwig Fuchs, who supported me with his deep expertise in both theory and practice. Raphael Neudert, who researched methods for log analysis with great zeal and scientific curiosity. And Tobias Reittinger, whose critical view has greatly enriched our joint work. I would also like to thank my co-researchers at Nexis and the university chair, who gave advice or challenged my work and thus made it better. Special mentions go to Dr. Michael Kunz, Dr. Matthias Hummer, Alexander Puchta, and Dr. Magdalena Glas. And finally, I would like to thank my loved ones: My parents, Susanne and Peter; my brother Michael, my sister-in-law Stephanie, and my niece Leonie; and my girlfriend Daniela. You always made me feel like I was on the right path. Thank you, sincerely.

Contents

Abstract	i
Acknowledgement	iii
List of Tables	vii
List of Figures	ix
List Of Abbreviations	xi
I Dissertation Outline	1
1 Introduction	3
2 Research Questions	5
3 Research Methodology	7
3.1 Guidelines for Design Science in Information Systems Research	8
3.2 Research Environment	9
4 Results	9
4.1 Foundations for Research Questions 1 and 2	11
4.2 Research Question 1: Quality Assessment	14
4.3 Research Question 2: Quality Maintenance	19
4.4 Research Question 3: Access Reviews	25
5 Conclusion and Future Work	29
References	30
II Research Papers	37
1 Optimization of Access Control Policies	39
2 Identity and Access Management Metrics	61
3 Transaction Logs in Access Control: Leveraging an Under-Utilized Data Source	99
4 Maintain High-Quality Access Control Policies: An Academic and Practice-Driven Approach.	113
5 A Framework for Managing Separation of Duty Policies	135
6 Monitoring Access Reviews by Crowd Labelling	147

7	Digital Nudges for Access Reviews: Guiding Deciders to Revoke Excessive Authorizations	165
---	---	-----

List of Tables

1	Research Questions and Subordinate Research Questions	7
2	Publications of this dissertation	11
3	Quality-related ACP properties identified in P1	12

List of Figures

1	Overview of publications and addressed SRQs	10
2	Relevant publications identified in the literature survey in P1	13
3	Strategic Identity and Access Management (IAM) goals derived from Oh and Pinsonneault [39], Hummer et al. [25] and Fuchs et al. [17] in P2. .	15
4	Method of demonstrating the analytical value of transaction logs for P3.	19
5	Schematic overview of the ACP maintenance framework proposed in P4.	21
6	Matrix visualization of example SoD classes and their pairwise mutual exclusions ("SoD matrix") as defined in P5.	23
7	Confusion matrix for authorizations as defined in P7.	28

List of Abbreviations

ACP	Access Control Policy
ABAC	Attribute-Based Access Control
CRA	Cyber Resilience Act
DORA	Digital Operational Resilience Act
DAC	Discretionary Access Control
IAM	Identity and Access Management
IF	Impact Factor
JSON	JavaScript Object Notation
LLM	Large Language Model
MAC	Mandatory Access Control
PGT	Pseudo Ground Truth
RQ	Research Question
RBAC	Role-Based Access Control
SAML	Security Assertion Markup Language
SCIM	System for Cross-Domain Identity Management
SOX	Sarbanes-Oxley Act
SMER	Static Mutually Exclusive Roles
SoD	Separation of Duty
SSoD	Static Separation of Duty
SRQ	Subordinate Research Question
WSC	Weighted Structural Complexity
XACML	eXtensible Access Control Markup Language

Part I

Dissertation Outline

1 Introduction

Access control is a basic requirement for protecting valuable resources. One of the oldest known access control mechanisms is a possibly 4,000-year-old lock that was found in the ruins of ancient Nineveh [13]. Like modern access control mechanisms, the lock was built to ensure that only those who are authorized can access a protected resource. Today, advancing digitalization confronts organizations with the modern counterpart of this old problem: Securing an ever-growing number of digital assets for a large number of users. Many organizations struggle to implement secure and efficient access controls, which often results in serious vulnerabilities. Numerous threat reports list broken or insufficient access control as one of the most frequent and critical IT security issues [2, 36, 6], potentially leading to data breaches, financial loss, or reputational damage. Recent trends, such as increasing regulatory requirements or non-human identities, further increase the technical and organizational complexity of this task. This dissertation contributes to addressing this challenge by examining how the quality of data, and in particular Access Control Policies (ACPs), can be assessed and improved to enable effective management of identities and their access.

This dissertation is positioned in information systems research. It presents research on Identity and Access Management (IAM) in the context of an organization. IAM is a domain of information security management driven by the needs for IT security, regulatory compliance and operational efficiency. It deals with provisioning users with secure and efficient access to digital resources. IAM employs a range of processes, policies and technologies which aim to ensure that IT security objectives are met without overly restricting users [17]. To ensure sufficient security, organizations define a security policy that specifies which access users are allowed or not allowed to make [44]. The security policy is subject to restrictions, such as the principle of least privilege, which defines that no user may receive more authorizations than he or she needs to fulfill his or her tasks. External regulations like the Sarbanes-Oxley Act (SOX) [47], the Basel accords [3], the Digital Operational Resilience Act (DORA) [15] or the Cyber Resilience Act (CRA) [14] also define restrictions that must be reflected in the security policy, such as the implementation of Separation of Duty (SoD), which aims to avoid conflicts of interest [20]. The authorizations defined in the security policy are enforced by an access control mechanism, which either permits or denies users access to resources. To make an access decision, the access control mechanism must evaluate ACPs, machine-processable rules that define authorizations and can be evaluated fully automated. The data structure of ACPs is defined by their access control model [42], with prominent examples such as Role-Based Access Control (RBAC)[43], Attribute-Based Access Control (ABAC)[24], Discretionary Access Control (DAC) or Mandatory Access Control (MAC)[33]. Several other entities are commonly processed as data in IAM. They are described in established standards such as OpenID Connect [16], Security Assertion Markup Language (SAML) [37], System for Cross-Domain Identity Management (SCIM) [26], or eXtensible Access Control Markup Language (XACML) [38]. Among the most commonly considered ones

are digital identities, i.e., digital representations of users in the context of an organization. Accounts, in contrast, are representations of users in the context of specific applications. Kunz et al. proposed a conceptual model that integrates central IAM entities with a unified terminology [31].

The quality of data is a critical success factor when working with information systems. While high data quality enables high-quality work results, low data quality leads to errors and acts as a cost driver [41, 22]. Data quality is described in terms of dimensions that reflect different aspects of the data's fitness for use [48]. Existing research has identified numerous data quality dimensions that are valid across application domains, such as accuracy, timeliness, or completeness. However, to tackle complex organizational problems, data quality must be conceptualized in the context of the data's use and composition [46]. In the context of IAM, the quality of ACPs and attributes is particularly important: ACPs are the *source of truth* for any access control mechanism. If they do not define authorizations accurately, users are granted excessive access, which causes a security vulnerability, or insufficient access, which causes business disruptions. In addition to ACP accuracy at a given time, maintainability-related quality criteria such as complexity and redundancy are critical for enabling policy engineers to retain accuracy at an acceptable cost [5]. Attributes are an abstract representation of entity properties that is often used to conceptualize data quality generically. In the context of IAM, it is well suited to describe quality-related properties of entities other than ACPs, e.g., digital identities or permissions. Furthermore, the accuracy of attributes can also be directly responsible for the correctness of authorization decisions, e.g., when evaluating ABAC policies.

The topic of this dissertation is the assessment and improvement of data quality in the context of IAM. It considers two data types: ACPs and attributes. Since ACPs are specific to the field of IAM, research on their quality is typically limited to this research realm. Moreover, it is often only considered as a secondary aspect of other contributions such as policy mining algorithms. This is evident in the lack of foundational research, such as the lack of consistent terminology and quality criteria (see publication P1). In contrast, attributes are a broadly applicable data type that is not bound to a particular information system domain. In fact, the structure of attributes in IAM does not differ fundamentally from that in other fields of application. This dissertation identifies a significant research gap in the quality of ACPs, which is thus the primarily focused data type. Attributes are being considered to describe properties of other IAM-related data. In addition to the generalizing consideration of data quality in IAM, this dissertation addresses the focus area of Access Reviews. Access reviews are an IAM process in which designated reviewers identify and correct errors in ACPs and related entities. This manual data verification process requires considerable effort, yet its effectiveness is limited [27]. To the best of the author's knowledge, existing research did not provide a generalized definition of Access Reviews at the time of starting this dissertation, and the assessment

and improvement of review decisions are only sparsely addressed. The dissertation also addresses this research gap.

The remainder of this dissertation is organized as follows. Chapter 2 defines the research questions addressed in the course of this work. Chapter 3 describes the methodological basis of this work. Chapter 4 presents research results in the form of published articles, and explains their contribution to the individual research questions. Chapter 5 concludes this work with an overview of the contributions and future research directions. The second part of this dissertation includes the full-texts of the published articles.

2 Research Questions

The structure of this dissertation is guided by three Research Questions (RQs), each addressing a key problem at a high level of abstraction. To explore these in greater depth, seven Subordinate Research Questions (SRQs) were formulated, each focusing on specific aspects relevant to answering the main RQs. These SRQs serve as a bridge to the individual contributions of the dissertation and define the scope of the research. The results presented in Chapter 4 contribute to the overarching RQs through their alignment with the corresponding SRQs.

A prerequisite for ensuring data quality is the evaluation of the quality at hand. Existing literature offers few definitions for criteria that determine the fitness for use of ACPs [5]. Most publications rely on proprietary working definitions and do not follow a coherent quality framework. As a result, a wide variety of quality criteria have been proposed, which are difficult to compare [30, 9]. Although the quality of attributes is repeatedly cited as a critical foundation, e.g., for modeling high-quality ACPs, there is little comprehensive insight into which entities and attributes are actually relevant in IAM [32]. The first RQ of this dissertation is thus:

RQ1: How to assess the quality of ACPs and attributes in IAM?

This dissertation addresses RQ1 by investigating which quality dimensions are relevant for ACPs and attributes in the context of IAM. SRQ 1.1 is thus defined as: *Which quality criteria are relevant for IAM data?* A specific challenge within this context is the identification of authorization errors, which are represented by the ACP quality dimension accuracy. Since the ground truth for authorizations as defined by an organization's security policy is usually not available in a structured form, their detection is error-prone and involves significant manual effort. Consequently, SRQ 1.2 is defined as: *How can excessive authorizations be identified?*

Based on the assessment of data quality, steps can be taken to improve the existing quality. Previous research has invested considerable effort into the initial modeling of high-quality access control policies. A wide range of techniques (such as policy mining, policy engineering, or hybrid approaches [18, 11]) have been proposed in the literature,

seeking to model ACPs while optimizing them according to selected quality criteria. However, the initial modeling of ACPs is a resource-intensive task that requires significant time and organizational effort. Once modeled, the quality of ACPs tends to deteriorate over time. This degradation can result from incorrect or suboptimal modifications, outdated rules due to changes in the organizational environment (e.g., shifts in employee responsibilities), or the overly permissive granting of new authorizations [49]. As a result, the security of the organization declines, the maintainability of ACPs decreases, and the initial investment in policy modeling is undermined. The quality of ACPs must thus be continuously improved in order to counteract this decay and maintain a sufficient level of security and operational efficiency. Unlike the initial modeling of ACPs, the improvement and maintenance of ACP quality have received little attention in academic research. Therefore, the second research question of this dissertation is:

RQ2: How to maintain high ACP quality continuously in an organization?

ACP maintenance within an organizational context is not purely an algorithmic problem. It must take into account IAM-specific processes, regulatory requirements, technologies, and data types. Relevant stakeholders must be identified, and domain experts should be involved in the maintenance process. As a result, both the technical and organizational complexity of ACP maintenance are considerable. To approach this problem step by step, SRQ 2.1 is defined as: *Which challenges are associated with ACP maintenance?* As a second step, existing solutions are reviewed in SRQ 2.2: *Which ACP maintenance approaches exist, and what are their limitations?* Building on this, conceptual frameworks are to be developed that define structures and responsibilities for the quality maintenance of ACPs and provide practical guidance for implementation. Accordingly, the third subordinate research question is SRQ 2.3: *How can ACP maintenance be implemented in the context of an organization?*

RQ1 and RQ2 aim to provide a comprehensive analysis of data quality in IAM. In addition, this dissertation delves into Access Reviews, an IAM process that serves as a concrete application case for both quality assessment and quality improvement. The primary goal of Access Reviews is to identify and resolve excessive authorizations. Secondary goals may also include identifying missing authorizations, inaccurate attribute values, or other error cases specific to an organization. During the Access Review process, a decision-maker (e.g., a department head or application administrator) reviews the authorizations of subjects under their responsibility. If errors in authorization assignments are found, they must be corrected by adjusting the relevant ACPs [27]. Besides reviewing effective permissions, Access Reviews can also involve the assessment of attributes or compositions of ACPs, such as description texts of ACPs or job function descriptions for employees. Therefore, Access Reviews are a specific IAM process aimed at detecting and correcting data errors. Access Reviews are often driven by regulatory requirements and require significant execution effort. Despite the considerable resources invested, the effectiveness of Access Reviews remains limited due to structural challenges (i.e., scale

and frequency, lack of knowledge, exceptional cases, and human errors [27]). Thus, the third research question of this dissertation is:

RQ3: How to enable effective Access Reviews?

The basis for evaluating the effectiveness of Access Reviews lies in determining the correctness of review decisions. To the best of the author's knowledge, there are no contributions in the existing scientific literature outside this dissertation addressing this issue. Therefore, SRQ 3.1 is formulated as: *How can the quality of Access Review decisions be determined ex-post?* Building on this, the next step is to investigate which factors positively or negatively influence the quality of Access Reviews. Since this is primarily a manually conducted process, the focus is on the work effectiveness of the reviewers. This leads to SRQ 3.2: *How can reviewers be guided to perform effective Access Reviews?* Table 1 summarizes all RQs and SRQs.

No.	Research Questions (RQ) and Subordinate Research Questions (SRQ)
RQ1	How to assess the quality of ACPs and attributes in IAM?
1.1	Which quality criteria are relevant for IAM data?
1.2	How can excessive authorizations be identified?
RQ2	How to maintain high ACP quality continuously in an organization?
2.1	Which challenges are associated with ACP maintenance?
2.2	Which ACP maintenance approaches exist, and what are their limitations?
2.3	How can ACP maintenance be implemented in the context of an organization?
RQ3	How to enable effective access reviews?
3.1	How can the quality of access review decisions be determined ex-post?
3.2	How can reviewers be guided to perform effective access reviews?

Table 1: Research Questions and Subordinate Research Questions

3 Research Methodology

This dissertation is positioned in the domain of information systems research. It addresses practice-relevant questions with scientific methods in order to generate insights for theory and practice. For this purpose, a methodology was used that supports the structuring and quality assurance of the research. The Design Science paradigm [23] provided fundamental orientation for this work. Additional research methodologies were applied as needed: The main contribution of publication P6 was developed using the Action Design Research method, a popular approach in Design Science research with particular focus on practical relevance and applicability [45]. The structured literature surveys presented in P1 and P2 followed the process model proposed by Levy and Ellis [34]. The semi-structured expert interviews presented in P4, P5, and P7 followed the guidelines defined by Adams [1]. The user study presented in P7 is analyzed using empirical methods.

3.1 Guidelines for Design Science in Information Systems Research

Hevner et al. propose seven guidelines for Design Science in Information Systems Research [23]. They aim to ensure a clear understanding, execution and evaluation of research. While these guidelines provide valuable orientation for structuring a research project, Hevner et al. argue that they should not be applied mandatorily or routinely. Instead, researchers must determine when, where and how the application of each guideline benefits a specific research project.

Guideline 1 - Design as an Artifact

Design Science research must produce a purposeful and viable artifact. The produced artifact must be described effectively to enable its implementation and application. The research results of this application provide IT artifacts in the form of frameworks and methods. They are presented and explained in the resulting publications.

Guideline 2 - Problem Relevance

The produced technology-based solution is relevant to business problems. It must thus add value for the practitioners that manage, design, implement, operate or evaluate information systems, or the technologies that enable their development and implementation. For this work, the problem relevance was ensured through the grounding in scientific literature and the exchange with IAM practitioners. Their feedback also encompassed design cycle iterations.

Guideline 3 - Design Evaluation

The utility, quality, and efficacy of an artifact must be rigorously demonstrated via well-executed evaluation methods. This includes observational, analytical, experimental, testing, or descriptive methods. The evaluation of the artifacts in P3, P4, P5, and P6 was carried out in large real-world organizations with IAM data used in productive environments.

Guideline 4 - Research Contributions

Design Science research must provide clear and verifiable contributions. The ultimate assessment for each research project is thus: "What are the new and interesting contributions?" This work offers clearly defined contributions to three main research questions by providing foundations or IT artifacts for data quality in IAM.

Guideline 5 - Research Rigor

Design Science relies on rigorous methods in the design and evaluation of artifacts. Researchers thus have to make good use of the existing knowledge base and use appropriate evaluation criteria. Still, Hevner et al. note that over-emphasis on rigorous methods can

lower the relevance of research and hinder its practical adaptation. It is thus necessary to assess the appropriateness of evaluation criteria, as the principal aim is to determine how well an artifact works rather than theorizing about it. In this work, rigor was assured through structured analysis of the knowledge base and adherence to the research methodologies as described.

Guideline 6 - Design as a Search Process

As Design Science is iterative, searching for an optimal design is typically not feasible. Instead, heuristic search strategies that produce a satisfyingly good design are preferable. The primary knowledge base of this dissertation was scientific literature. To supplement it, the author sought exchange with IAM practitioners throughout the research. In some cases, industry partners also contributed directly to the research results: Contributions of the research papers P4, P5, and P7 were developed in semi-structured expert interviews with IAM practitioners.

Guideline 7 - Communication of Research

Design Science research must be effectively presented to both technology-oriented and management-oriented audiences. This required a balance between presenting enough detail to enable an implementation of the produced artifact, and describing its use in an organizational context. The primary medium to feed the results of this dissertation into the knowledge base are the scientific publications. In addition, the results were also made accessible to the participating industry partners.

3.2 Research Environment

This dissertation was written in a hybrid research setting in cooperation with the University of Regensburg and the company Nexis. Both organizations employ experts in IT security and IAM, who were available for exchange of knowledge, advice, and research collaborations. The University of Regensburg provided an academic environment: in addition to supervising this dissertation, a doctoral seminar was offered there, in which plans and results of research could be presented and discussed. Nexis offered an industrial environment with access to IAM practitioners in medium- and large-sized organizations. Parts of the research were carried out within the DEVISE research project. The project was funded by the German Federal Ministry of Education and Research (BMBF). It influenced the scope of this dissertation and contributed to its financing.

4 Results

The research of this dissertation resulted in seven scientific publications P1 – P7. Each of them contributes to answering at least one SRQ raised in Chapter 2. Their index numbers 1 – 7 were chosen to simplify the overview of the addressed SRQs. At the

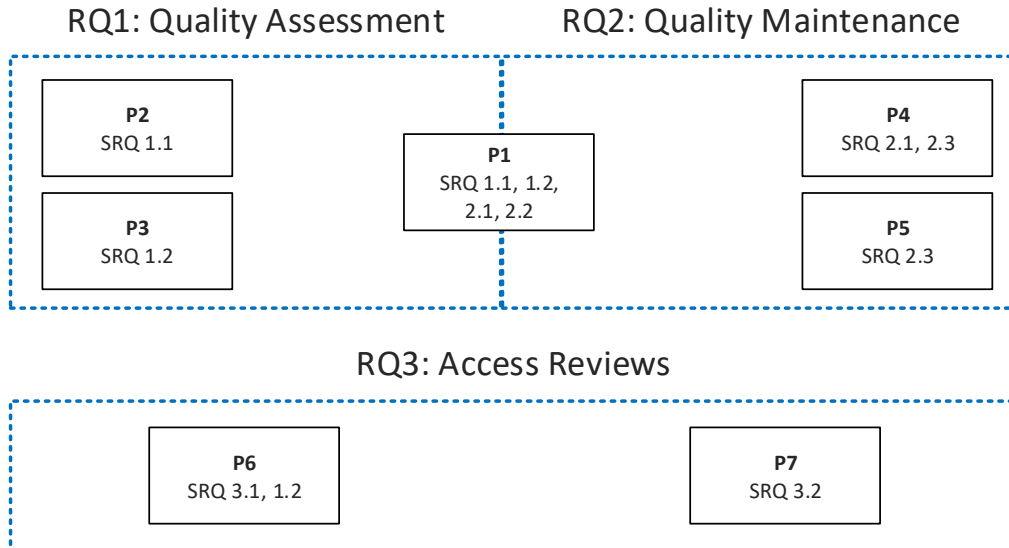


Figure 1: Overview of publications and addressed SRQs

time of submitting this dissertation, six publications (P1, P3, P4, P5, P6, and P7) were successfully published. The remaining publication P2 is currently under review. The seven publications contribute to SRQs as follows:

Publication P1 lays the foundation for RQ1 and RQ2 by structuring and analyzing the state of the art in existing scientific literature. It provides contributions to SRQs 1.1 and 1.2 as well as 2.1 and 2.2.

P2 and P3 contribute to RQ1 by addressing the SRQs 1.1 and 1.2 respectively: P2 contributes to SRQ 1.1 by identifying and analyzing quality criteria for digital identities and ACPs. P3 introduces transaction logs as a possible data source for determining authorization accuracy and thus contributes to SRQ 1.2.

Both P4 and P5 contribute to RQ2: P4 identifies five challenges for ACP maintenance based on expert interviews and existing literature, contributing to SRQ 2.1. It then proposes a framework for ACP maintenance that contributes to SRQ 2.3. P5 proposes a framework for managing SoD policies and thus also contributes to SRQ 2.3.

Finally, P6 and P7 contribute to RQ3. P6 defines the task of identifying Access Review errors as an instance of crowd labelling. It proposes a method to identify possible errors, which contributes to SRQ 3.1. As the identified errors reveal excessive authorizations, P6 also contributes to SRQ 1.2. P7 presents expert interviews and a user study to analyze the effect of digital nudges on the work of Access Review deciders (i.e., reviewers), contributing to SRQ 3.2.

Figure 1 gives a graphical overview of all publications and addressed SRQs. Table 2 summarizes all publications with their full citation and the CORE ranking or Impact Factor (IF) of the conferences or journals they were presented in. In the remainder of this chapter, the seven publications are presented in detail.

	Publication	Ranking
P1	Sascha Kern, Thomas Baumer, Sebastian Groll, Ludwig Fuchs, and Günther Pernul. Optimization of access control policies. <i>Journal of Information Security and Applications</i> , 70:103301, 2022.	3.8 IF
P2	Thomas Baumer, Sascha Kern, Ludwig Fuchs, and Günther Pernul. Identity and Access Management Metrics. Submitted to: <i>ACM Computing Surveys</i> .	-
P3	Sascha Kern, Thomas Baumer, Raphael Neudert, and Günther Pernul. Transaction Logs in Access Control: Leveraging an Under-Utilized Data Source. In <i>IFIP Annual Conference on Data and Applications Security and Privacy</i> , pages 413-424. Springer, 2025.	B CORE
P4	Sascha Kern, Thomas Baumer, Ludwig Fuchs, and Günther Pernul. Maintain high-quality access control policies: an academic and practice-driven approach. In <i>IFIP Annual Conference on Data and Applications Security and Privacy</i> , pages 223-242. Cham: Springer Nature Switzerland, 2023.	B CORE
P5	Sebastian Groll, Sascha Kern, Ludwig Fuchs, and Günther Pernul. A framework for managing separation of duty policies. In <i>Proceedings of the 19th International Conference on Availability, Reliability and Security</i> . ACM, 2024.	B CORE
P6	Sebastian Groll, Sascha Kern, Ludwig Fuchs, and Günther Pernul. Monitoring access reviews by crowd labelling. In Simone Fischer-Hübner, Costas Lambrinouidakis, Gabriele Kotsis, A. Min Tjoa, and Ismail Khalil, editors, <i>Trust, Privacy and Security in Digital Business</i> , pages 3-17, Cham, 2021. Springer International Publishing.	B CORE
P7	Thomas Baumer, Tobias Reittinger, Sascha Kern, and Günther Pernul. Digital nudges for access reviews: guiding deciders to revoke excessive authorizations. In <i>Twentieth Symposium on Usable Privacy and Security</i> , pages 239-258, 2024.	B CORE

Table 2: Publications of this dissertation

4.1 Foundations for Research Questions 1 and 2

Publication P1 lays the groundwork for addressing RQ1 and RQ2. This chapter presents its contributions, which define central concepts and the research background for following results.

P1. Optimization of Access Control Policies

P1 analyzes the state of research on ACP quality assessment. For the first contribution, it identifies 16 properties of ACPs that are related to their quality, and summarizes them at a high level of abstraction. Each of these properties was used in existing literature to

No.	Property	Optimum	Affects evaluation
C1	Accuracy	Max	Yes
C2	Excessive authorizations	Min	Yes
C3	Missing authorizations	Min	Yes
C4	Maintainability	Max	No
C5	Understandability	Max	No
C6	Sem. meaningfulness	Max	No
C7	Complexity	Min	No
C8	Redundancy	Min	No
C9	Conflicts	Min	Yes
C10	Grade of automation	Max	Yes
C11	Evaluation runtime	Min	Yes
C12	Similarity to optimal state	Max	No
C13	Risk	Min	No
C14	Completeness	Max	Yes
C15	Usage	-	No
C16	Relevance	-	No

Table 3: Quality-related ACP properties identified in P1

determine aspects of ACPs' fitness for use. Out of the 16 identified quality-related ACP properties, 14 (C1 - C14) are considered quality dimensions. They have an unambiguous optimum, so the quality of ACPs can be improved by either minimizing or maximizing these properties. The remaining two properties, C15 and C16, were repeatedly used in literature to determine ACP quality, but are not direct indicators of ACPs' fitness for use. For each of the 16 properties, approaches for quantification exist in the literature. However, the measurement of these properties exceeded the scope of P1. While seven identified properties impact policy evaluation by the access control mechanism, the remaining nine describe the ACPs at rest, and do not influence their evaluation. Table 3 summarizes the identified properties. P1 contributes to answering SRQ 1.1 as it shows which quality criteria exist for ACPs and how they are relevant for their fitness for use. To the best of the author's knowledge, this represents the most comprehensive compilation of high level ACP quality dimensions at the time of submitting this dissertation. Table 3 summarizes all 16 properties.

The principal contribution of P1 is a structured literature review, carried out in accordance with the guidelines proposed by Levy and Ellis [34]. Its literature scope is set for scientific publications that propose methods for optimization of role-based or attribute-based ACPs in relation to at least one defined ACP quality dimension. To define this scope more precisely, a total of six quality dimensions were selected for research: *C2 Excessive Authorizations*, *C3 Missing Authorizations*, *C8 Redundancy*, *C9 Conflicts*,

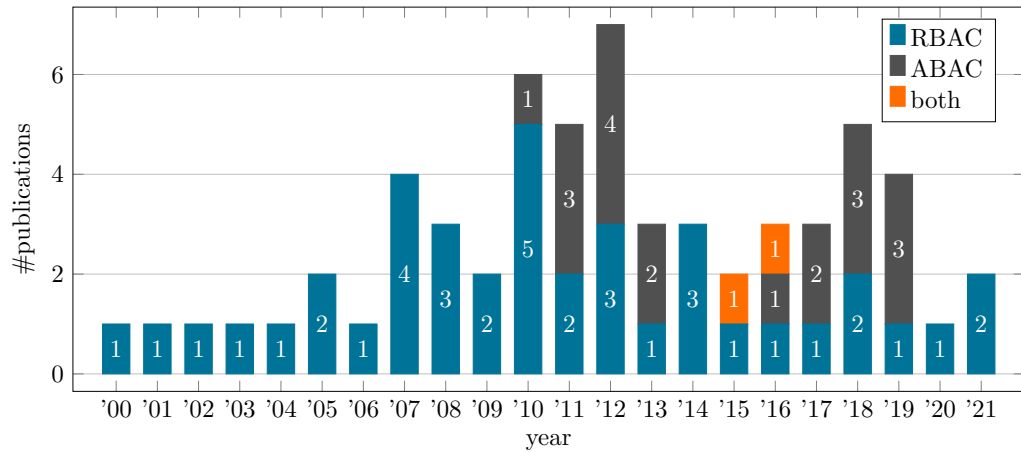


Figure 2: Relevant publications identified in the literature survey in P1

C7 Complexity and *C10 Grade of Automation*. The selection of these quality dimension was based on two main reasons: (i) Beckerle and Martucci [5] conducted semi-structured expert interviews and a literature analysis to identify critical requirements for obtaining usable ACPs. They also developed metrics for their quantification and conducted two user studies to evaluate them. Based on their results, they argue that the main aim of ACP optimization should be to improve accuracy and maintainability of ACPs. They identify six crucial optimization criteria that can be mapped exactly to the six identified quality dimensions: 1. "Allow no more than the owner wants to be allowed", 2. "Allow everything the owner wants to be allowed", 3. "A rule must not be fully covered by another rule of the same rule set", 4. "Two rules belonging to the same rule set must not conflict", 5. "Minimize the number of rule set elements" and 6. "Minimize maintenance effort in a changing system". To the best of the author's knowledge, it is the only scientific publication that documents a structured research process for the development of ACP optimization objectives, and provides a conclusive evaluation. (ii) Each of these six quality dimensions is frequently used in literature to determine ACPs' fitness for use.

The literature research and filtering process yielded 61 publications that fit within the scope of the survey, published in the years 2000 to 2021. 42 of them address the improvement of roles, 21 the improvement of ABAC policies, and two address both. The literature was categorized according to three criteria: (i) The optimization objective, i.e. which quality dimension was improved by a proposed approach. (ii) The type of research artifact. Most publications presented at least one of the following four types of research artifact designed to improve ACP quality: A process model, an algorithm, a tool, or an ACM extension. (iii) Data usage: Many optimization approaches rely on one or more of the following data sources: Entity attributes, access logs, or transaction logs. The obtained literature catalogue was analyzed and classified with regard to each optimization objective, as well as the considered optimization scenarios. P1 thus contributes to SRQ 2.2. Figure 2 summarizes the identified publications, ordered chronologically by their year of publication, and the addressed ACMs.

In the third part, P1 analyzes the literature body for further insights. As a contribution

to SRQ 2.1, it analyzes the limitations and maintenance challenges addressed in the obtained literature catalogue. One challenge is obtaining all the context data required for maintenance approaches. P1 argues that the availability of access logs is limited, which means that approaches relying on access logs have limited applicability in practice. The second challenge is integrating maintenance algorithms into existing business processes. A crucial factor is whether an approach can generate optimization recommendations that can be individually accepted or rejected. While some approaches generate single, independent ACP update steps, others create a full sequence of interdependent updates that lead to a new, improved ACP state. Their ability to produce update operations that can be decided individually thus affects the applicability of ACP maintenance approaches due to mandatory approval processes and other business constraints. Thirdly, P1 analyzes how the concept of minimal perturbation is considered. ACP maintenance faces the challenge of updating existing policies with minimal changes to minimize update effort while retaining most of the existing policy structure. Finally, P1 contributes to SRQ 1.2 by identifying and discussing three prototypical approaches for identifying excessive authorizations from the literature. Manual identification is the de facto practice standard, but limited in its effectiveness. Identification based on access logs is established in literature, but significantly limited in practice due to log availability. Identification based on transaction logs is proposed by some authors, but not yet established in theory or practice.

Contribution of P1

P1 identifies 16 quality-related ACP properties (SRQ 1.1) and analyzes prototypical approaches to identify authorization inaccuracies (SRQ 1.2). A structured literature survey analyzes approaches for ACP maintenance (SRQ 2.2) and a discussion highlights overarching challenges for these approaches (SRQ 2.1).

4.2 Research Question 1: Quality Assessment

Publications P2 and P3 focus on the assessment of data quality. P2 identifies quality metrics within IAM, including data quality. P3 introduces transaction logs as a data type for ACP quality assessment and investigates how they can be used to assess a ground truth for authorization accuracy.

P2. Identity and Access Management Metrics

P2 researches metrics that are relevant for strategic goals of IAM. This includes, but is not limited to, central data quality metrics for digital identities and ACPs, contributing to SRQ 1.1. As a foundation for this work, the central goals of IAM were derived from the scientific literature. Subsequently, a structured literature review was conducted using the process model proposed by Levy and Ellis [34] to search for metrics measuring the achievement of strategic IAM goals. Through an iterative process, IAM metrics were identified and synthesized from the literature. Subsequently, the impact of the metrics

on the achievement of the respective goals was examined, as well as the stakeholders addressed by each metric. Finally, the metrics and their implications were presented in a structured manner. The discussion highlights key findings from the research process.




Figure 3: Strategic IAM goals derived from Oh and Pinsonneault [39], Hummer et al. [25] and Fuchs et al. [17] in P2.

The primary strategic goal of IAM is cost reduction. This is addressed through related IAM goals, including ensuring security, regulatory compliance, and operational efficiency. These cost-related goals are complemented by quality-related goals that enable effective IAM. They include sufficient quality of digital identities and policies, as well as IAM processes, and IAM software [39, 25, 17]. Together, these cost-related and quality-related goals define the scope of the literature survey. Figure 3 summarizes the strategic IAM goals in scope. Publication P2 adopts the term "metric" following the definitions of Black et al. [7] and Chew et al. [10]: A metric is an abstract, somewhat subjective attribute, e.g., the complexity of a password. A measure is a concrete, objective attribute, e.g., the length or entropy of a password. Analysts can collect various measures to quantify metrics. Measures themselves can be calculated using different formulas. P2 thus aims to compile abstract metrics that represent the effectiveness of activities and the goal achievement in IAM at a conceptual level. The concrete implementation of the metrics via specific measures and formulas is beyond scope of this work.

The literature survey identified seven IAM perspectives that are represented by metrics: The quality and life cycles of digital identities and ACPs, IAM process quality, IAM software quality, and a management and governance perspective. ACPs have been introduced earlier in this dissertation. Digital identities are a representation of an individual in a digital context. According to Pfitzmann and Hansen [40], they are defined as a subset of attributes which identifies an individual among a set of individuals. The life cycles of ACPs and digital identities describe the processing of these data from their creation up to their deactivation. IAM employs a range of common processes and digital technologies, such as joiner, mover, and leaver processes, or access control systems, directories and provisioning systems. Finally, the governance perspective focuses on metrics that are important for the steering of IAM. These seven perspectives are divided into two areas: The identity life cycle, policy life cycle, and the management and governance perspective are assigned to the first area, as they primarily address the goal of cost reduction. The perspectives of ACP and digital identity quality, IAM processes and IAM software are assigned to the second area, which focuses primarily on quality-related metrics. The seven perspectives cover a total of 43 metrics with a unanimous optimization direction (minimization or maximization).

The collected metrics are analyzed in terms of their practical application scenarios.

The key findings are presented in P2. In addition to the textual presentation, the metrics are categorized: First, the positive and negative influencing factors of the properties observed by each metric on the strategic IAM goals are analyzed. A categorization summarizes whether a high or low metric value indicates a positive or negative influence on the respective IAM goal. Second, the target audiences of the metrics are analyzed. Five typical audiences are identified: (i) A governance audience covers management, compliance and audit responsibilities. (ii) Identity administrators are typically supervisors or human resource employees. (iii) Policy administrators design and maintain ACPs. (iv) Operations are responsible for keeping IAM systems running, including roles such as help desk, DevOps or provisioning engineers. (v) Users rely on IAM to access digital resources. The metrics are categorized according to their intended target audience. Finally, a discussion highlights the insights gained from the work and gives advice on utilizing the metrics.

At the time of submitting this dissertation, P2 is under review at the Journal *ACM Computing Surveys*. It was submitted in August 2023.

Contribution of P2

P2 identifies IAM metrics and analyzes their audience and relationship to strategic IAM goals. This includes quality metrics for digital identities and ACPs, as well as for their life cycles, contributing to SRQ 1.1.

P3.Transaction Logs in Access Control: Leveraging an Under-Utilized Data Source

P3 contributes to SRQ 1.2 by introducing a type of context data in IAM research that may help determine inaccuracies in ACPs and attributes. The contribution consists of three parts: (i) P3 introduces and formalizes the concept of transaction logs for access control. (ii) Based on scientific literature and industry standards, it analyzes how transaction logs can be collected within a representative IAM infrastructure. (iii) P3 anchors the analysis of transaction logs within IAM processes. P3 evaluates the contributions through a case study using real-world enterprise data.

The existing literature commonly considers four types of data relevant for ACP maintenance: The access control matrix, attributes, access logs, and transaction logs. The access control matrix displays the effective authorizations resulting from existing ACPs. It is required to determine the authorization state *as-is* and ensure that authorizations remain accurate after modifying policies, e.g., to reduce their complexity. Attributes of entities (such as users or permissions) provide valuable context information and can be used to align ACPs to real-world concepts, that is, provide semantic meaning. However, neither the access control matrix nor attributes provide a ground truth to determine authorization inaccuracies in an existing ACP set. Access logs are records of historic authorization invocations. They are commonly used in the literature to identify inaccurate authorizations. However, their usefulness is constrained by factors such as availability, detectable inaccuracy types, interpretive ambiguity, and regulatory restrictions, as they

may enable workplace surveillance. Transaction logs are frequently cited in the literature (although with inconsistent terminology such as *update logs*, *event logs* or *change histories*). While several publications argued for their analytical value, to the best of the author’s knowledge, no previous scientific work has formally defined them or proposed methods for their use in identifying inaccurate authorizations.

```

1  {
2    "transactionLog": [ {
3      "operation": "create",
4      "entity": "John Doe",
5      "entityType": "identity",
6      "timestamp": "2025-02-03T18:48:50Z"
7    }, {
8      "operation": "create",
9      "entity": "Reviewer",
10     "entityType": "role",
11     "timestamp": "2025-02-04T18:48:50Z"
12    }, {
13     "operation": "update",
14     "entity": "John Doe",
15     "entityType": "identity",
16     "attribute": "isExpert",
17     "attributeValue": "true",
18     "timestamp": "2025-02-05T18:48:50Z"
19    }, {
20     "operation": "grant",
21     "identity": "John Doe",
22     "role": "Reviewer",
23     "timestamp": "2025-02-06T18:48:50Z"
24    } ]
25  }

```

Listing I.1: Example for transaction logs showing user “John Doe” becoming a reviewer.

While this listing represents them in JavaScript Object Notation (JSON), transaction logs are not bound to a specific data format.

P3 defines transaction logs in IAM as structured change histories for relevant IAM data. Listing I.1 shows an example of four transaction log entries. Despite inconsistent terminology, the concept of transaction logs is well-established in other fields of IT, such as database revision management [12], software version control systems [8], and digital ledger technologies [4]. The publication provides a formal definition for transaction logs for IAM using first-order predicate logic. It then introduces a schematic centralized IAM architecture based on scientific literature and industry standards, and analyzes where transaction logs can be obtained. With this foundation, P3 proposes mapping transaction log entries to process instances to gain meaningful insights into ACPs and possible inaccuracies. It discusses mapping log entries to three representative IAM processes: The Joiner process, the Mover process, and the Access Review process. Transaction log events recorded during the Joiner process document the creation of a digital identity with all of its attributes and the initial assignment of the authorizations required by a user. They show an attribute and authorization state that is more likely to be correct than the current state data, as it is timely and no entitlement accumulation has taken

place yet. Mover process logs document data changes that occurred during a change of a user's affiliation within an organization, e.g., a department, function, or job title. They can show meaningful relationships between data change events, e.g. a user receiving an authorization shortly after joining a department. Mover process logs can thus be used to analyze why individual users received or lost specific authorizations, or to derive strong correlations between the assignment of and revocation of attributes and authorizations in large data volumes. Access Review logs document the revocation of authorizations after a human reviewer decided that they are excessive. Due to the manual inspection and the tendency of reviewers to accept existing authorizations in case of doubt (see P7), authorization revocations from Access Reviews are very likely to show an actual over-authorization, which can be used to identify further ones (see P6). Transaction logs thus allow for a meaningful representation of these real-world IAM processes. P3 argues that all three processes provide a structure for transaction log analysis that can be used to assess a ground truth for data accuracy.

The contributions are validated in a case study with real-world enterprise data provided by a large pharmaceutical technology producer. The data set comprises 35,559 digital identities that are associated with 253,253 user accounts, and 420,387 permissions that are managed via 13,801 roles. The company operates two dedicated IAM systems: One is used for data provisioning, and the other for data analysis and governance tasks. Both systems provide an integrated view on the relevant IAM data (i.e., the digital identities, permissions, user accounts, roles, proprietary authorization rules, and supplementary data such as department structures). Both systems log changes in data they observe to comply with regulations and provide error support, and thus provide a comprehensive transaction log as defined in P3. The authors implemented a data processing pipeline that collected and normalized selected transaction logs and generated association rules using the Apriori algorithm. By processing log events assignable to Mover process instances, the authors identified association rules that showed strong correlations between attribute changes and role assignments or unassignments. Such rules can be used to identify contradictions in the current-state IAM data, highlighting possibly inaccuracies in authorizations or attribute values. After an initial validation against known automation rules, selected rules were presented to IAM experts from the partnering company, who confirmed the plausibility of selected rules and emphasized their value for authorization analysis. The authors identified 1,157 user-role assignments which contradicted 20 association rules with a high confidence score (≥ 0.95). These cases were considered suspicious because a user with a given set of attributes was unlikely to hold the given role. Through selective review, the company experts confirmed that these included likely cases of excess authorizations that warranted further investigation. However, since this would have required extensive effort on the part of the partners, they were unable to provide a conclusive assessment of the extent of excessive assignments actually found. Figure 4 summarizes the validation approach. Due to page restrictions of the conference proceedings, the validation could not be presented in the submitted research paper. However, a research demo containing

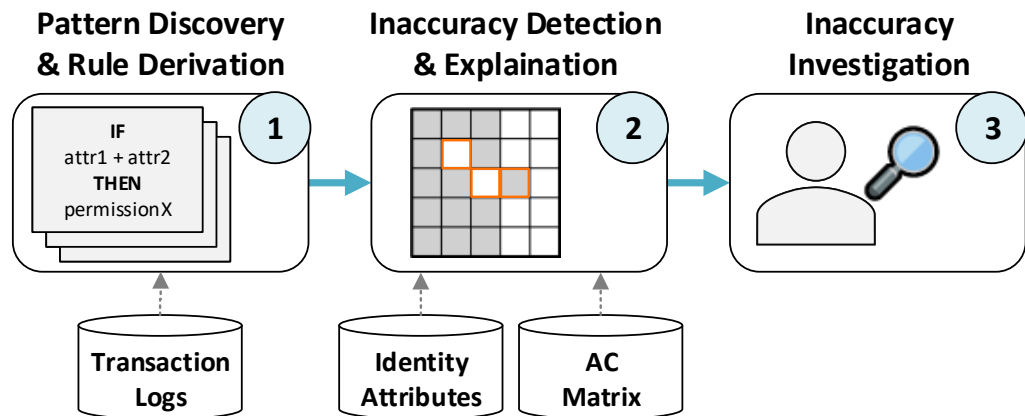


Figure 4: Method of demonstrating the analytical value of transaction logs for P3.

the underlying implementation and a sample data set were made available at GitHub¹. In summary, the authors found two suitable sources for transaction logs in a large enterprise, thus validating contributions (i) and (ii). By linking log events with joiner process instances, relevant insights into real authorization structures could be gained, validating contribution (iii). P3 thus established a foundation for the analysis of transaction logs. It provides groundwork for identifying inaccurate authorizations and attributes without access logs.

Contribution of P3

P3 introduces the use of transaction logs to analyze IAM data. It establishes a foundation for identifying inaccurate authorizations and attributes and thus contributes to SRQ 1.2.

4.3 Research Question 2: Quality Maintenance

Publications P4 and P5 examine how to continuously improve the quality of ACPs in an organization. They address technical and organizational challenges related to ACP maintenance: P4 proposes a framework that reduces organizational complexity by defining responsibilities and tasks for semi-automatized ACP maintenance. P5 addresses the more specific case of SoD policies and formalizes the process of quality maintenance using SoD matrices.

P4. Maintain High-Quality Access Control Policies: An Academic and Practice-Driven Approach

Publication P4 provides three contributions to RQ2: First, it identifies and summarizes five major challenges to ACP maintenance based on scientific literature and interviews with real-world IAM experts, contributing to SRQ 2.1. Second, it proposes a framework that provides a structure for the maintenance of ACP, contributing to SRQ 2.3. Third,

¹<https://github.com/TransactionLogs/Availability>

the proposed framework is evaluated with a real-world IAM data set of a large financial services company.

A structured literature search yielded 17 publications that describe problems related to ACP maintenance. This foundation was expanded with six semi-structured expert interviews in accordance with the methodology of Adams [1]. This research yielded numerous specific problems and illustrative examples, which were summarized into five overarching challenges. Examples of each of these challenges were described in both the literature examined and the expert interviews. The five key challenges identified are: (i) The amount and complexity of policies, (ii) distributed knowledge, (iii) the importance of undisturbed business processes, (iv) organizational and regulatory restrictions, and (v) attribute quality. The identified challenges are described in detail in P4.

Based on this problem analysis, this publication proposes a framework for the maintenance of ACPs. The development followed the Design Science paradigm. The framework defines activities related to ACP maintenance and locates them in four IAM domains: (i) Governance, (ii) the IAM team, (iii) IT & domain experts, and (iv) the maintenance environment. The governance domain defines strategic goals and reviews their achievement. The IAM team has the operational responsibility for the maintenance of ACPs. They define ACP quality objectives based on the strategic IAM goals, and review their achievement. They are responsible for implementing the ACP maintenance process and aiding the stakeholders involved in it. The IT & domain experts include people with operational knowledge who make the actual maintenance decisions. Their involvement is necessary because ACP updates require human decisions (challenge iv), and the knowledge required for these decisions lies distributed over an organization (challenge ii): For example, a department head might know which tasks their employees have to carry out, or a system administrator might know which effects certain permissions have in their systems. The IAM team members typically have a good understanding of the overall IAM architecture and policy structure of an organization, but cannot answer such detailed questions themselves. The amount and complexity of the policies (challenge i) also makes it preferable for the IAM team to out-source the maintenance of ACPs to the business domains which are affected by them. The maintenance environment, which is the fourth framework domain, provides tools and applications which automate and aid the ACP maintenance. It continuously evaluates the ACP state to identify quality problems and generates the necessary maintenance actions. For example, a presumed excessive authorization might trigger a permission withdrawal, or the detection of an SoD conflict might require a resolution. If necessary, it prompts these maintenance actions as recommendations to the responsible IT & domain experts, who may approve or reject them. Maintenance actions that are accepted, or do not require human approval, are applied to improve ACP quality. It is also possible to just inform stakeholders about possible quality problems if no unambiguous maintenance action can be derived, e.g. if a policy requires a periodic check for timeliness. Figure 5 summarizes the proposed framework.

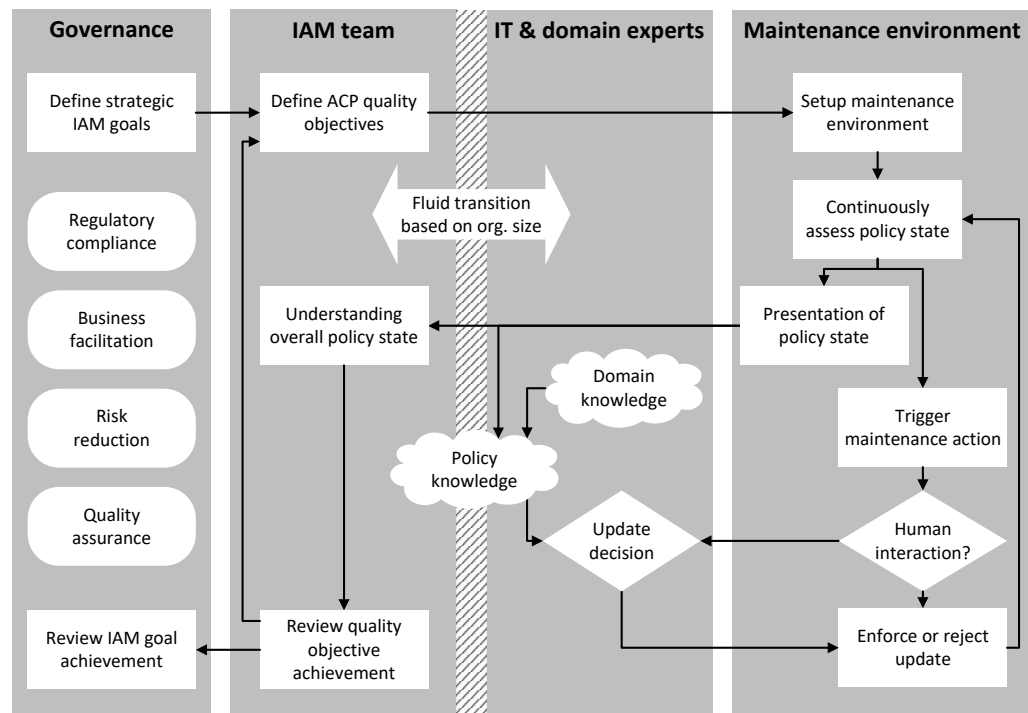


Figure 5: Schematic overview of the ACP maintenance framework proposed in P4.

A case study was conducted to evaluate the proposed framework. For this purpose, the authors cooperated with IAM practitioners of a large financial service provider company. These research partners provided the authors with read-only access to the productively used IAM data within the company's IT infrastructure. The company's IAM was based on RBAC with 2,385 *Business Roles* that managed a total of 290,914 permissions for 7,344 users within 662 application systems. The framework was instantiated in the organizational context of the company. The research partners obtained the stakeholder function of IAM governance, while the authors took the role of the IAM team. Since no real domain experts of the company were accessible, this role was approximated by programmatically rejecting 20% of proposed maintenance actions. In their role of IAM governance, the research partners formulated strategic IAM goals: To improve IAM operations, the data overview should be improved and unnecessary complexity reduced. To reduce risk, their goals also included removing excessive authorizations and improving master data quality. Given limited resources and environmental restrictions, the authors defined two maintenance objectives for the scope of the evaluation: Reduce complexity and reduce redundancy of the ACPs. These quality metrics were measured using the Weighted Structural Complexity (WSC) (as defined in [50]) and the ratio of redundantly assigned authorizations. The authors then implemented the maintenance environment with ACP analysis and maintenance capabilities within the company's IT infrastructure. It included capabilities to collect, normalize, browse and visualize the available IAM data, calculate quality metrics, and generate possible maintenance actions to achieve a quality improvement. The authors simulated maintenance execution within two improvement cycles, which led to the removal of 22.39% of redundant authorizations and a reduction

in WSC of 9.56%. The research partners confirmed that the achieved improvements in quality were considerable, reducing data complexity and improving general overview. The maintenance environment implementations and simulated ACP updates were handed over to the research partners for practical application.

Contribution of P4

P4 contributes to SRQ 2.1 by identifying five core challenges in ACP maintenance. It then proposes and evaluates a framework for ACP maintenance which addresses SRQ 2.3.

P5. A Framework for Managing Separation of Duty Policies

Publication P5 focuses on maintaining high-quality SoD policies, a specific type of ACP that aims to prevent conflicts of interest. P5 provides three contributions that address SRQ 2.3: (i) Seven semi-structured expert interviews with SoD managers from large, compliance-driven companies were conducted. The aim was to understand common procedures for SoD management and to provide a foundation for the development of the framework. (ii) A framework was developed that outlines stakeholders and responsibilities to manage SoD policies in a maintainable way. It focuses on making SoD policies easily understandable by enriching them semantically. (iii) The conceptual framework is formalized for RBAC and evaluated using a real-world IAM data set of a large financial services company.

The expert interviews were conducted in natural language with seven SoD managers from six companies, following the interview methodology by Adams [1]. The interviews followed a prepared questionnaire. The authors deviated from the questionnaire to follow up on ambiguities or other relevant points that arose during the interviews. Each interview was structured into three blocks: The first block aimed to introduce the interviewees to the topic and collect general information like their job position, the company's size, and their motivation for managing SoD policies. In the second block, the authors examined how SoD policies are structured in the interviewees' companies. It revealed that SoD policies are typically managed in one of three forms: (i) Pairwise mutually exclusive permissions are the simplest and most granular way to define SoD conflicts. They are easily understood by humans, but quickly result in high data complexity. (ii) Pairwise mutually exclusive roles work the same way, but on a higher level of aggregation, as they define exclusions between roles. (iii) In five of the six organizations, the bulk of SoD exclusions are defined via SoD classes. An SoD class defines a semantic concept, e.g., "market" or "market follow-up" (a back-office position that can approve loans). SoD classes are attributed to permissions: In this example, a permission that allows granting a loan would be attributed the class "market", and a permission that allows to approve a loan "market follow-up". By defining pairwise exclusions between SoD classes, SoD managers implicitly also define exclusions between permissions and roles, but are able to manage SoD exclusions on a high level of abstraction. Figure 6 shows a matrix representation of

example SoD classes and exclusions, commonly referred to as "SoD matrix". Although common in practice, to the best of the authors' knowledge, the concept of SoD classes had not yet been introduced to scientific literature and formalized prior to this publication. In the third interview block, the interviewees were asked about their companies' SoD life cycles and processes. The aim was to identify relevant stakeholders, and to determine how SoD policies are created and maintained in practice. A summary of the interview results is presented in the publication.

	Market	Market Follow-Up	Audit	Risk Controlling	Accounting	Legal	Compliance	Trade	Payment Traffic	Fund Mgt.
Market										
Market Follow-Up										
Audit										
Risk Controlling										
Accounting										
Legal										
Compliance										
Trade										
Payment Traffic										
Fund Mgt.										

Figure 6: Matrix visualization of example SoD classes and their pairwise mutual exclusions ("SoD matrix") as defined in P5.

Based on the interview results, a framework for SoD policy management was developed following the Design Science paradigm. It defines stakeholders, activities, and three prototypical structures for SoD policies, which are designed to support maintainability of SoD policies. Three types of SoD policies are defined: The SoD matrix, pairwise mutually exclusive permissions and pairwise mutually exclusive roles. The SoD matrix is a representation of SoD classes and their exclusions, as identifiers in the expert interviews. It offers a visual representation of the managed SoD exclusions, which can be maintained without domain knowledge of individual permissions or users. The pairwise mutually exclusive permission and role policies are a specific form of Static Separation of Duty (SSoD) and Static Mutually Exclusive Roles (SMER) policies as defined in scientific literature [35]. The three SoD policy specifications are rooted in the integrated IAM data model by Kunz et al. [32]. They are formalized in RBAC, demonstrating that they are compatible with the ACM and can be transformed into classic SMER constraints. The framework defines three stakeholders and their activities for SoD maintenance: (i) IAM Governance analyzes regulatory requirements and risks. They define and maintain the SoD matrix. (ii) The IAM Team is responsible for IAM project management, and thus for the technical SoD implementation. They organize the definition of SoD policies and ensure their enforcement. (iii) Domain Experts have the domain knowledge required for fine-grained maintenance of SoD policies. They are responsible for assigning SoD classes

to individual permissions, and for the maintenance of fine-grained, mutually exclusive role or permission policies.

The framework was evaluated in cooperation with an SoD manager from a large financial services provider, who had also participated in the expert interviews. For the scope of the evaluation, the authors were provided access to the productively used IAM data set of the company. The conceptual overview of stakeholders and activities was generally agreed upon as it originated from the expert interviews. Thus, two evaluation criteria were defined: (i) **Technical applicability:** The evaluation should confirm that the SoD matrix and pairwise mutual exclusions are compatible with existing SoD specifications, and that they can be enforced as such. Since the transformation of pairwise mutually exclusive roles and permissions into SMER constraints is trivial, the transformation of SoD classes into SMER constraints was focused. (ii) **Simplified policy management:** The evaluation should confirm that the proposed framework provides a complexity reduction in comparison to established SMER constraints. The company employed around 4,500 employees. The data set contained around 10,000 digital identities, and an RBAC model, with 2,494 roles and 7,972 permissions. It defined 14 SoD classes with 32 exclusions among these classes. The SoD classes were assigned to 209 roles and 274 permissions. Two algorithms were implemented for the evaluation, which transformed the SoD classes and their constraints along with the assigned roles and permissions into classic SMER constraints. They were applied on the evaluation data to generate 12,295 pairwise SMER constraints. The evaluation partner confirmed their correctness. The results showed that: (i) The technical applicability criterion is satisfied. A real-world SoD matrix can be represented by SMER constraints and enforced as such. This means that the proposed framework can be applied as a layer of abstraction for policy management, while the existing SMER specifications from literature remain valid. (ii) The simplified policy management criterion is satisfied. The SoD matrix required a total of 529 data entities that need to be maintained, whereas the equivalent SMER constraints comprised 12,295 entities. The complexity of the SoD matrix is thus far below the complexity of the corresponding SMER constraints. While it would be possible to reduce the complexity of SMER constraints algorithmically, this would also make them harder to understand, as the semantic meaning of the pairwise exclusions would be lost. An algorithmic complexity reduction of SMER constraints would therefore merely exchange one aspect of maintainability for another (see P1).

Contribution of P5

P5 presents findings from expert interviews on the management of SoD policies. It proposes a framework and RBAC-based formalizations for the creation and management of easily maintainable SoD policies, which contribute to SRQ 2.3.

4.4 Research Question 3: Access Reviews

Publications P6 and P7 aim to improve the effectiveness of Access Reviews. P6 proposes an approach for determining the decision quality of reviewers, while P7 explores the use of digital nudges to increase the number of detected inaccuracies.

P6. Monitoring Access Reviews by Crowd Labelling

Publication P6 is the first publication of this dissertation to address Access Reviews as a specific instance of IAM data quality. The research process followed the Action Design Research method by Sein et al. [45]. This research method is positioned in the Design Science research domain and aims to develop and evaluate IT artifacts in an organizational setting. It defines a research process with four stages: The problem formulation stage introduces a practical problem and conceptualizes a research opportunity. The building, intervention, and evaluation stage generates a design for an IT artifact in a cyclic development process. The reflection and learning stage broadens the conceptualized solution of a particular problem to a class of problems. Finally, the formalization and learning stage codifies the generalized result and reflects on insights gained during the research process. The research question formulated in P6 is: *"How can low-quality review decisions in an Access Review be identified ex-post?"* and thus contributes to answering SRQ 3.1. P6 provides four contributions: (i) It introduces an abstract definition of Access Reviews based on the conceptual IAM data model proposed by Kunz et al [32]. (ii) P6 shows that Access Review decision quality is an instance of the crowd labelling quality research problem. (iii) A crowd labelling approach that identifies low-quality labelling decisions is adapted for Access Reviews, enabling identification of low-quality Access Review decisions. (iv) The approach is evaluated using Access Review decision data of a real-world enterprise.

P6 formalizes Access Reviews based on the conceptual IAM data model by Kunz et al. The formalization is based on four types of entities, each with a set of attributes: Employees, accounts, business roles, and permissions. The entities are connected via relations in an entity-relation graph. The effective authorizations are defined by the permissions inherited by employees in the graph. P6 defines two classes of Access Reviews: (i) In class 1 reviews, entity attributes are reviewed and, if inaccurate, corrected. (ii) In class 2 reviews, the relations between entities are reviewed. Relations are removed if they violate the organizational security policy. To narrow the research scope, P6 focuses exclusively on class 2 reviews in the building, intervention and evaluation stage.

Crowd sourcing is a collaborative activity in which tasks are distributed to workers who execute them in a decentralized manner [29]. Like Access Reviews, crowd sourcing faces the challenge of determining the quality of results: Different workers can produce results of very different quality, and typically no structured ground truth is available to evaluate the correctness of all decisions. Crowd labelling is a subdomain of crowd sourcing. It is typically used to create labels for large data sets, e.g., for machine learning. A class 2 review decision is defined as a crowd labelling task: A decision d

of reviewer r concerns the relation between two entities e_1 and e_2 . It comprises a set of features from the two entities, and the information on whether the two entities should remain linked. The features x_{e1} and x_{e2} are the attributes of entities e_1 and e_2 . The review decision is represented as a binary label y : While $y = true$ represents the "keep relation" decision, $y = false$ represents "remove relation". A review decision can thus be expressed as $d_r(e_1, e_2) = \{x_{e1}, x_{e2}, y\}$. An Access Review with n decisions is defined as a set $AR = \{d_i | i \in 1, \dots, n\}$. Since excessive authorizations are of particular interest for Access Reviews, P6 formulates the research problem: "How to assess the quality of each Access Review decision d_i with the label $y = true$ ". Following these definitions, Access Reviews are an instance of crowd labelling, with three specific characteristics: (i) While typical crowd labelling decisions may involve more than one possible label, Access Review decisions consist only of binary labels. (ii) While typical crowd labelling tasks involve multiple workers labelling the same data, any Access Review decision regarding two entities is made by a single reviewer. (iii) The quality of positive Access Review decisions is considered more important than the quality of negative Access Review decisions.

Based on these formalizations, P6 adapts a quality assessment approach for crowd labelling decisions to evaluate Access Review decisions. The approach adapted from Geva and Saar-Teschansky [19] is designed for crowd labelling decisions that are created by exactly one reviewer. It estimates decision quality by generating a Pseudo Ground Truth (PGT), which is treated as the correct decision. The PGT is then compared with the actual decisions of workers to estimate the quality of their decision, and to create a quality score for each worker. Since each labelling decision is made by a single worker, it is not possible to estimate the PGT by comparing the results of other workers for the same decision. Instead, a base model is generated for every worker. A base model predicts a label y for any given feature vector $\{x_{e1}, x_{e2}\}$ based on all past decisions made by the actual worker. The assessed quality score of a worker is thus computed by measuring the accordance of their labels with the labels predicted by the base models trained on the decisions of other workers. Any decision which falls below a defined quality threshold t is assumed to be inaccurate.

The adapted approach was evaluated on a real-world Access Review data set of a large telecommunications provider. It comprised the results of one Access Review campaign regarding account-permission relations, summing up to 71,464 decisions concerning 10,891 employees and 1,623 permissions. The adaptations showed three challenges that made it difficult to identify incorrect labels with $y = true$: (i) 93.9% of the decisions were positive, which created a strong bias towards positive decisions in the base models. (ii) Many base models had to make decisions for feature vectors which did not share any attributes with those used during their training. (iii) Additional noise arose from base models that were trained on very limited data, in some cases even a single decision. These issues were tackled by introducing an *experience score* which determined the qualification of base models for any given review decision. This experience score was

used to weight predicted labels during quality score calculation, which significantly increased the number of presumed labelling errors identified. The approach was evaluated based on a subset of the review data. It yielded 33 potentially inaccurate reviewer decisions, of which 30 (90.9%) were confirmed to be actual review errors by IAM experts of the telecommunications company. P6 thus contributes to SRQ 3.1 by proposing a method to assess the quality of Access Review decisions. In the chosen evaluation scenario, the proposed approach de facto identified excessive authorizations based on Access Review decisions. P6 thus also contributes to SRQ 1.2, demonstrating the potential of Access Review decision data for identifying possible authorization errors that were not found by reviewers. Furthermore, the quality measuring approach by Geva and Saar-Teschansky was improved through further development for the subclass of Class 2 Access Reviews.

Contribution of P6

P6 contributes to SRQ 3.1 by adapting a crowd labelling approach to assess the decision quality of Access Reviews. As it identifies authorizations that are likely excessive, this approach also contributes to SRQ 1.2.

P7. Digital Nudges for Access Reviews: Guiding Deciders to Revoke Excessive Authorizations

P7 contributes to SRQ 3.2 by examining how the effectiveness of reviewers is influenced by digital nudges. Thus, this publication focuses on the human component of the reviews. The contribution is threefold: (i) P7 formalizes the Access Review problem. (ii) Ten expert interviews were conducted to analyze how five core challenges of Access Reviews as defined by Jaferian et al. [27] are affected by 13 digital nudges established in scientific literature. (iii) A user study with 102 participants analyzes how the *Choice Defaults* nudge [28] influences reviewer behaviour.

To define the Access Review problem, P7 introduces a confusion matrix for authorizations. It compares the correct authorizations (Positive P and Negative N), as defined by an authorization's security policy, with the actually granted authorizations (Predicted Positive PP and Predicted Negative PN), as defined by an access control matrix. The effectively granted authorizations are thus defined as $PP = \text{True Positive TP} + \text{False Positive FP}$, i.e. the sum of correctly and incorrectly granted authorizations. Withheld authorizations are consequently defined as $PN = \text{False Negative FN} + \text{True Negative TN}$, i.e. the sum of incorrectly and correctly withheld authorizations. The sensitivity $SEN = \frac{TP}{P}$ thus represents the rate of correctly granted authorizations. The specificity $SPC = \frac{TN}{N}$ is the rate of correctly withheld authorizations, yielding a balanced accuracy $BA = \frac{SEN+SPC}{2}$ of 100% if the access control matrix exactly reflects the authorizations specified in the security policy. Given that Access Reviews typically aim to revoke excessive authorizations, P7 defines the Access Review problem as the task of designing an Access Review in such a way that decision makers can reduce the False Discovery

Rate $FDR = \frac{FP}{PP}$ without lowering BA. This formalization allows an exact definition of the effectiveness of an Access Review and forms the basis for conducting statistical hypothesis testing in user studies. Figure 7 shows the defined confusion matrix.

		Authorization	
		Positive PP	Negative PN
Security Policy	Positive P	TP	FN
	Negative N	FP	TN

Figure 7: Confusion matrix for authorizations as defined in P7.

Ten semi-structured expert interviews were conducted with practitioners that have 5 - 19 years of IAM experience. The interviews followed the guidelines proposed by Adams [1]. They aimed to assess the general potential of digital nudges for Access Reviews. The interviews started with an introduction covering the experts' background and experience and an explanation of the interview approach. The interviewees were then asked for a qualitative assessment of the effect of digital nudges on Access Review challenges. During the interviews, 13 digital nudges established in scientific literature were mapped to the five Access Review challenges raised by Jaferian et al. For each of the 65 resulting pairings, interviewees were asked to provide qualitative examples of how a given nudge would affect a given Access Review challenge, and to rate the presumed positive or negative impact on the challenge quantitatively on a five-point Likert scale. The interviews provided a qualitative knowledge base that helps assess possible positive and negative impacts of the surveyed digital nudges on the effectiveness of Access Reviews. The qualitative assessments were aggregated to offer an overview of the perceived nudge effectiveness.

P7 then studies the effects of the *Choice Defaults* nudge on Access Reviews in a user study with 102 participants. The Choice Defaults nudge was selected because it is widely regarded as one of the most effective nudges in scientific literature, and because the expert interviews revealed divergent assessments of its possible effects on Access Reviews. Study participants had to perform an Access Review on a synthetic data set using a review tool. They had to review 160 authorizations and decide whether they should remain *granted* or be *revoked*. The study authors designed the data set to be easily understood while having an unambiguous ground truth. The study participants were split into three groups of 34 participants each: A group with the Choice Defaults nudge suggesting a *grant* decision for all reviewed authorizations, a group with a *revoke* decision default, and a control group without any Choice Defaults nudge applied. All study participants received a single-page instruction in natural language designed to provide a sufficient foundation for informed review decisions without revealing the correct answers. At the end of the review, each study participant was asked to complete a questionnaire to assess the perceived workload using the NASA Task Load Index [21]. The study was designed to take approximately 20-30 minutes per participant.

The result data shows that almost all study participants improved the Access Review problem. While the required time differed substantially, it was not significantly correlated

with the result quality. The reviewers were inclined to accept authorizations, although the study data was balanced and the introduction did not give any other indication. The applied Choice Defaults nudge did not significantly affect the result quality, but reduced the required time and the perceived stress. The configuration with a Default Reject decision also achieved a reduction of the bias towards accepting authorizations. The participants' self-assessment correlated with their actual performance, suggesting that reviewers have a realistic perception of their review quality. This highlights the importance for reviewer motivation for the effectiveness of Access Reviews. The study also identified a subgroup of reviewers making obviously flawed decisions, referred to as "spammers". These individuals represent an easily identifiable source of errors. Overall, P7 showed that digital nudges have potential for improving the effectiveness of Access Reviews, without undermining reviewer autonomy.

Contribution of P7

P7 contributes to SRQ 3.2 by studying how 13 digital nudges could potentially influence the effectiveness of human deciders in Access Reviews. A user study with 102 participants analyzes the impact and potential benefits of the *Choice Defaults* nudge in greater detail.

5 Conclusion and Future Work

The research presented in this dissertation provides contributions to the assessment and improvement of data quality in the context of IAM. It considers the data types of ACPs and attributes, as well as the focus area of Access Reviews. Four publications analyze the assessment of data quality in response to the first research question: Publications P1 and P2 define generic quality criteria for ACPs and digital identities. Publications P1, P3 and P6 contribute to assessing the accuracy of ACPs. After P1 identifies existing approaches from the literature, P3 formally introduces transaction logs to IAM and analyzes possible uses of this data type for the detection of ACP inaccuracies. P6 proposes a method that identifies possible inaccuracies using Access Review decisions. Three publications present results on data quality maintenance, addressing the second research question: Publications P1 and P4 work out key challenges in a structured literature survey and semi-structured expert interviews. P1 provides a structured overview of existing approaches for ACP quality improvement and analyzes their limitations. On this basis, P4 and P5 propose frameworks that aid in the continuous maintenance of high-quality ACPs. Finally, publications P6 and P7 address the focus area of Access Reviews, contributing to the third research question: P6 formally defines Access Reviews as an instance of crowd labelling, and proposes a method to determine the quality of reviewer decisions. P7 formalizes the Access Review problem for an empirical analysis and investigates the effects and potential use of digital nudges on reviewer effectiveness.

This work offers several directions for future research. First, scientific literature lacks

foundational work on ACP quality: A unified conceptualization of quality dimensions that is valid across different access control models, and an analysis of their inter-dependencies or their organizational impact, do not yet exist. Second, the automated assessment of ACP accuracy remains a fundamental problem: While existing literature offers many approaches to create or update policies to reflect a given target authorization state, little research aims to determine the correctness of the authorization state itself. Approaches that do typically rely on access logs, a data type with significant limitations in availability and analytical power (see publication P3). Approaches that analyze transaction logs, e.g. by using machine learning techniques, may provide a valuable alternative. Third, integrating algorithms for quality improvement into existing IAM operations is a challenge hardly acknowledged in scientific literature: Many optimization approaches assume that an existing policy set can be simply re-generated from scratch, but do not acknowledge the need for human-understandable and individually decidable update steps (see publication P1). The disproportionately higher addressing of policy mining in scientific literature underlines this blind spot. Approaches that create meaningful policy update recommendations can address this research gap. Fourth, Access Reviews are a widely established IAM process that binds significant resources while remaining limited in its effectiveness. Approaches to improve the effectiveness of Access Reviews, e.g. through a targeted use of digital nudges, would have high application potential. At the same time, Access Review results represent a source of identified authorization errors. These can be used to identify further errors (see P6) and thus to address the aforementioned lack of ground truth. Fifth, recent advancements on Large Language Models (LLMs) promise new ways to interpret data with semantic context. To the best of the author's knowledge, existing research has not yet attempted to use them in the considered scenarios. LLM-based approaches might serve to identify IAM data errors, or provide data overview for reviewers in ways that current approaches cannot.

Bibliography

- [1] William C Adams. Conducting semi-structured interviews. *Handbook of practical program evaluation*, pages 492–505, 2015.
- [2] Andrew Van der Stock, Brian Glas, Neil Smithline, and Torsten Gigler. Owasp top 10:2021. Technical report, Open Web Application Security Project (OWASP), 2021. Accessed: 2025-06-14. URL: <https://owasp.org/www-project-top-ten/>.
- [3] Basel Committee on Banking Supervision. Principles for effective risk data aggregation and risk reporting (bcbs 239), 2013. Accessed: 2025-06-14. URL: <https://www.bis.org/publ/bcbs239.pdf>.
- [4] Roman Beck, Jacob Stenum Czepluch, Nikolaj Lollike, and Simon Malone. Blockchain—the gateway to trust-free cryptographic transactions. In *Twenty-Fourth European Conference on Information Systems (ECIS), İstanbul, Turkey, 2016*, pages 1–14. Springer Publishing Company, 2016.
- [5] Matthias Beckerle and Leonardo A Martucci. Formal definitions for usable access control rule sets from goals to metrics. In *Proceedings of the ninth symposium on usable privacy and security*, pages 1–11, 2013.
- [6] Beyond Identity. Former employees admit to using continued account access to harm previous employers, Feb 2022. Accessed: 2025-06-14. URL: <https://www.beyondidentity.com/blog/great-resignation-impact-on-company-security>.
- [7] Paul E Black, Karen Scarfone, Murugiah Souppaya, et al. Cyber security metrics and measures. *Wiley Handbook of Science and Technology for Homeland Security*, pages 1–15, 2008.
- [8] Kevin Buffardi. Assessing individual contributions to software engineering projects with git logs and user stories. In *Proceedings of the 51st ACM technical symposium on computer science education*, pages 650–656, 2020.
- [9] Manuel Cheminod, Luca Durante, Fulvio Valenza, and Adriano Valenzano. Toward attribute-based access control policy in industrial networked systems. In *2018 14th IEEE international workshop on factory communication systems (WFCS)*, pages 1–9. IEEE, 2018.

- [10] Elizabeth Chew, Marianne Swanson, Kevin Stine, Nadya Bartol, Anthony Brown, and Will Robinson. Performance measurement guide for information security. Special Publication (NIST SP) 800-55 Revision 1, National Institute of Standards and Technology, Gaithersburg, MD, 2008. Accessed: 2025-06-14. URL: https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=152183.
- [11] Carlos Cotrini, Thilo Weghorn, and David Basin. Mining abac rules from sparse logs. In *2018 IEEE European Symposium on Security and Privacy (EuroS&P)*, pages 31–46. IEEE, 2018.
- [12] Tony Davis, Gail Shaw, and Kalen Delaney. *SQL Server Transaction Log Management*. Simple Talk Pub., 2012.
- [13] Encyclopaedia Britannica, Inc. Lock. Encyclopaedia Britannica, 1998. Accessed: 2025-06-14. URL: <https://www.britannica.com/technology/lock-security>.
- [14] European Commission. Proposal for a regulation on horizontal cybersecurity requirements for products with digital elements (cyber resilience act), 2022. Accessed: 2025-06-14. URL: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=COM:2022:454:FIN>.
- [15] European Parliament and Council. Regulation (eu) 2022/2554 on digital operational resilience for the financial sector (dora), 2022. Accessed: 2025-06-14. URL: <https://eur-lex.europa.eu/eli/reg/2022/2554/oj>.
- [16] OpenID Foundation. Openid connect core 1.0, 2014. Accessed: 2025-06-14. URL: https://openid.net/specs/openid-connect-core-1_0.html.
- [17] Ludwig Fuchs and Gunther Pernul. Supporting compliant and secure user handling—a structured approach for in-house identity management. In *The Second International Conference on Availability, Reliability and Security (ARES'07)*, pages 374–384. IEEE, 2007.
- [18] Ludwig Fuchs and Günther Pernul. Hydro–hybrid development of roles. In *Information Systems Security: 4th International Conference, ICISS 2008, Hyderabad, India, December 16-20, 2008. Proceedings 4*, pages 287–302. Springer, 2008.
- [19] Tomer Geva and Maytal Saar-Tsechansky. Who’s a good decision maker? data-driven expert worker ranking under unobservable quality. 2016.
- [20] Sebastian Groll, Ludwig Fuchs, and Günther Pernul. Separation of duty in information security. *ACM Computing Surveys*, 2025.
- [21] Sandra G Hart. Nasa-task load index (nasa-tlx); 20 years later. In *Proceedings of the human factors and ergonomics society annual meeting*, volume 50, pages 904–908. Sage publications Sage CA: Los Angeles, CA, 2006.

- [22] Bernd Heinrich, Marcus Hopf, Daniel Lohninger, Alexander Schiller, and Michael Szubartowicz. Data quality in recommender systems: the impact of completeness of item content data on prediction accuracy of recommender systems. *Electronic Markets*, 31:389–409, 2021.
- [23] Alan Hevner, Salvatore T. March, Jinsoo Park, and Sudha Ram. Design Science in Information Systems Research. *Management Information Systems Quarterly*, 28(1):75–105, 2004.
- [24] Vincent C Hu, David Ferraiolo, Rick Kuhn, Arthur R Friedman, Alan J Lang, Margaret M Cogdell, Adam Schnitzer, Kenneth Sandlin, Robert Miller, Karen Scarfone, et al. Guide to attribute based access control (abac) definition and considerations (draft). *NIST special publication*, 800(162):1–54, 2013.
- [25] Matthias Hummer, Sebastian Groll, Michael Kunz, Ludwig Fuchs, and Günther Pernul. Measuring identity and access management performance-an expert survey on possible performance indicators. In *ICISSP*, pages 233–240, 2018.
- [26] Internet Engineering Task Force (IETF). System for cross-domain identity management: Core schema (rfc 7643), 2015. Accessed: 2025-06-14. URL: <https://tools.ietf.org/html/rfc7643>.
- [27] Pooya Jaferian, Hootan Rashtian, and Konstantin Beznosov. To authorize or not authorize: helping users review access policies in organizations. In *10th Symposium On Usable Privacy and Security (SOUPS 2014)*, pages 301–320, 2014.
- [28] Mathias Jesse and Dietmar Jannach. Digital nudging with recommender systems: Survey and future directions. *Computers in Human Behavior Reports*, 3:100052, 2021.
- [29] Aniket Kittur, Jeffrey V Nickerson, Michael Bernstein, Elizabeth Gerber, Aaron Shaw, John Zimmerman, Matt Lease, and John Horton. The future of crowd work. In *Proceedings of the 2013 conference on Computer supported cooperative work*, pages 1301–1318, 2013.
- [30] Michael Kunz, Ludwig Fuchs, Michael Netter, and Günther Pernul. How to discover high-quality roles? a survey and dependency analysis of quality criteria in role mining. In *International Conference on Information Systems Security and Privacy*, pages 49–67. Springer, 2015.
- [31] Michael Kunz, Alexander Puchta, Sebastian Groll, Ludwig Fuchs, and Günther Pernul. Attribute quality management for dynamic identity and access management. *Journal of information security and applications*, 44:64–79, 2019.
- [32] Michael Kunz, Alexander Puchta, Sebastian Groll, Ludwig Fuchs, and Günther Pernul. Attribute quality management for dynamic identity and access management.

- Journal of Information Security and Applications*, 44:64–79, 2019. URL: <https://www.sciencedirect.com/science/article/pii/S2214212618301467>, doi:10.1016/j.jisa.2018.11.004.
- [33] Donald C Latham. Department of defense trusted computer system evaluation criteria. *Department of Defense*, 198, 1986.
- [34] Yair Levy and Timothy J Ellis. A systems approach to conduct an effective literature review in support of information systems research. *Informing Science*, 9, 2006.
- [35] Ninghui Li, Mahesh V Tripunitara, and Ziad Bizri. On mutually exclusive roles and separation-of-duty. *ACM Transactions on Information and System Security (TISSEC)*, 10(2):5–es, 2007.
- [36] Microsoft Corporation. Microsoft digital defense report 2024, 2024. Accessed: 2025-06-14. URL: <https://cdn-dynmedia-1.microsoft.com/is/content/microsoftcorp/microsoft/final/en-us/microsoft-brand/documents/Microsoft%20Digital%20Defense%20Report%202024%20%281%29.pdf>.
- [37] OASIS. Assertions and protocols for the oasis security assertion markup language (saml) v2.0, 2005. Accessed: 2025-06-14. URL: <https://docs.oasis-open.org/security/saml/v2.0/saml-core-2.0-os.pdf>.
- [38] OASIS. extensible access control markup language (xacml) version 3.0, 2013. Accessed: 2025-06-14. URL: <https://docs.oasis-open.org/xacml/3.0/xacml-3.0-core-spec-os-en.html>.
- [39] Wonseok Oh and Alain Pinsonneault. On the assessment of the strategic value of information technologies: conceptual and analytical approaches. *MIS quarterly*, pages 239–265, 2007.
- [40] Andreas Pfitzmann and Marit Hansen. Anonymity, unlinkability, unobservability, pseudonymity, and identity management-a consolidated proposal for terminology. 2005.
- [41] Thomas C Redman. The impact of poor data quality on the typical enterprise. *Communications of the ACM*, 41(2):79–82, 1998.
- [42] Pierangela Samarati and Sabrina Capitani De Vimercati. Access control: Policies, models, and mechanisms. In *International school on foundations of security analysis and design*, pages 137–196. Springer, 2000.
- [43] Ravi S Sandhu. Role-based access control. In *Advances in computers*, volume 46, pages 237–286. Elsevier, 1998.
- [44] Ravi S Sandhu and Pierangela Samarati. Access control: principle and practice. *IEEE communications magazine*, 32(9):40–48, 1994.

- [45] Maung K Sein, Ola Henfridsson, Sandeep Puroo, Matti Rossi, and Rikard Lindgren. Action design research. *MIS quarterly*, pages 37–56, 2011.
- [46] Diane M Strong, Yang W Lee, and Richard Y Wang. Data quality in context. *Communications of the ACM*, 40(5):103–110, 1997.
- [47] U.S. Congress. Sarbanes-oxley act of 2002, public law 107-204, 2002. Accessed: 2025-06-14. URL: <https://www.govinfo.gov/content/pkg/PLAW-107publ204/pdf/PLAW-107publ204.pdf>.
- [48] Richard Y Wang and Diane M Strong. Beyond accuracy: What data quality means to data consumers. *Journal of management information systems*, 12(4):5–33, 1996.
- [49] Tianyin Xu, Han Min Naing, Le Lu, and Yuanyuan Zhou. How do system administrators resolve access-denied issues in the real world? In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pages 348–361, 2017.
- [50] Zhongyuan Xu. *Mining Meaningful Role-Based and Attribute-Based Access Control Policies*. Ph.d. dissertation, Stony Brook University, Stony Brook, NY, 2014. Accessed: 2025-06-14. URL: <https://www.proquest.com/openview/5a7ebae3b544417cc9ce2802a4b47031/1?pq-origsite=gscholar&cbl=18750>.

Part II

Research Papers

1 Optimization of Access Control Policies


Current status:	Published
Journal:	Journal of Information Security and Applications (JISA)
Date of acceptance:	July 31, 2022
Full citation:	Sascha Kern, Thomas Baumer, Sebastian Groll, Ludwig Fuchs, and Günther Pernul. Optimization of access control policies. <i>Journal of Information Security and Applications</i> , 70:103301, 2022.
Authors contributions:	Sascha Kern 45% Thomas Baumer 30% Sebastian Groll 10% Ludwig Fuchs 5% Günther Pernul 10%

Journal Description: Journal of Information Security and Applications (JISA) focuses on the original research and practice-driven applications with relevance to information security and applications. JISA provides a common linkage between a vibrant scientific and research community and industry professionals by offering a clear view on modern problems and challenges in information security, as well as identifying promising scientific and "best-practice" solutions. JISA issues offer a balance between original research work and innovative industrial approaches by internationally renowned information security experts and researchers.



Contents lists available at ScienceDirect

Journal of Information Security and Applications

journal homepage: www.elsevier.com/locate/jisa

Optimization of Access Control Policies

Sascha Kern^{a,*}, Thomas Baumer^a, Sebastian Groll^{a,b}, Ludwig Fuchs^a, Günther Pernul^b^a Nexis GmbH, Franz-Mayer-Straße 1, Regensburg, 93053, Bavaria, Germany^b University of Regensburg, Universitätsstraße 31, Regensburg, 93053, Bavaria, Germany

ARTICLE INFO

Keywords:

Access Management
Data quality
Policy optimization
Policy maintenance
Role-Based Access Control
Attribute-Based Access Control

ABSTRACT

Organizations undertake complex and costly projects to model high-quality Access Control Policies (ACPs). Once built, these policies must be maintained and managed in an ongoing process to keep their quality high. Insufficient maintenance leads to inaccurate authorization decisions and increases the policies' administrative effort and susceptibility to errors. While the initial modeling of ACPs has received significant research interest, their optimization is not yet covered as broadly. This work provides a theoretical foundation for ACP quality and its optimization. Furthermore, it analyzes how existing research addresses optimization of ACPs with regard to six crucial optimization dimensions. It presents a structured literature survey tracing these optimization dimensions, the contributed research artifact and data requirements. Building on this literature catalogue, this work elaborates on inaccuracies for user permission assignments, data availability, minimal perturbation and recommendation-based optimization.

1. Introduction

The organizational structures and IT infrastructures of modern companies are subject to constant change. Routine operations like departmental changes of employees, changing responsibilities or the integration of new application systems into the IT landscape require an adaptation of IT security configurations. This includes updating Access Control Policies (ACPs), machine-processable rules that define authorizations and can be evaluated in a fully automated manner to determine which accesses an employee is allowed to make [1]. Due to changing environmental conditions, ACPs that were once of high quality lose accuracy over time [2,3]. Moreover, ACPs proliferate over time, as policy administrators may over-grant access to conform with immediate business needs [4,5] or update policies in an erroneous or non-optimal way. Besides hard errors, ACP proliferation leads to lower comprehensibility and maintainability, which increases the ACPs' administrative cost and their proneness for further errors [6,7].

Incorrect or overly permissive access decisions leave companies vulnerable to insider threats. A malicious or careless insider can harm an organization severely, with consequences spanning from unintentional incidents to sabotage, fraud or espionage [8]. In contrast, overly restrictive access decisions prevent employees from doing their work, leading to costly interruptions in operations and task backlogs. Recent

studies estimate that the average annual cost of insider threats for companies reach \$11.45 million in 2020 [8,9]. The implementation of effective Identity and Access Management (IAM) measures, which follow the principle of least privilege [10], is hence mandated by major regulatory frameworks and IT security standards.¹

Maintenance measures, which optimize the quality of ACPs in a continuous manner, are a fundamental requirement for providing an accurate level of security with reasonable administrative effort over a longer period of time, and for maintaining the investment made in the initial modeling of high quality ACPs [11–13]. As the initial modeling of ACP sets with high quality requires high time and financial effort [14], maintenance processes aim to improve the quality of an ACP set by applying updates that leave the existing state intact. ACP maintenance is commonly approached in two types of processes: First, access reviews are a process where responsible humans (such as a department head) review ACPs for entities in their responsibility (for example roles assigned to their employees) and try to find and rectify inaccuracies. The effectiveness of access reviews is limited, since reviewers have to check large amounts of data in a largely manual process and have limited information to make a qualified decision [14–16]. Second, ACP refinement processes aim to improve the quality of an ACP set by updating it in a (semi-)automated manner without

* Corresponding author.

E-mail address: sascha.kern@wiwi.uni-regensburg.de (S. Kern).URLs: <https://nexis-secure.com/> (L. Fuchs), <https://www.uni-regensburg.de/informatik-data-science/wi-pernul/startseite/index.html> (G. Pernul).

¹ Such as the Sarbanes–Oxley Act (One Hundred Seventh Congress of the United States of America, 2002), the Basel Accords (Basel Committee on Banking Supervision), the European General Data Privacy Regulation (The European Parliament and the Council of the European Union, 2016), the ISO 27000 standards (International Organization for Standardization, 2013), or the BSI Grundschutz (Bundesamt für Sicherheit in der Informationstechnik, 2019).

<https://doi.org/10.1016/j.jisa.2022.103301>

Available online 15 September 2022

2214-2126/© 2022 Published by Elsevier Ltd.

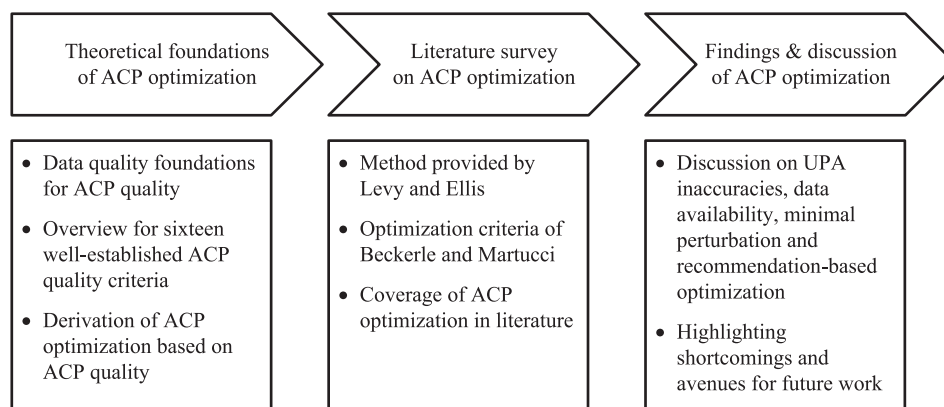


Fig. 1. General approach of the study.

deconstructing it [2]. Both approaches will be considered in the course of this work.

This work contributes to an optimization of ACPs by addressing the following **research question**: *Which methods for the optimization of roles and Attribute-Based Access Control (ABAC) policies are present, and what are their advantages and limitations?*. The contribution of this work is threefold: (i) We provide a definition of ACP quality as an instance of the data quality concept and supplement it with a collection of 16 ACP properties that are repeatedly used in literature to determine ACP quality. Building on this, we define ACP optimization as an improvement of the quality of existing ACPs. (ii) We conduct a structured literature survey based on the methodology proposed by Levy and Ellis [17]. Our scope is set for scientific publications describing means for optimization of ACPs that satisfy at least one of the six optimization objectives of Beckerle and Martucci [7]. After obtaining a literature catalogue for publications on optimization of ACPs, we categorize and analyze the literature catalogue. (iii) Finally, we build on findings from the literature survey and discuss important aspects of ACP optimization in more detail. At first, we discuss prototypical approaches to identify User Permission Assignment (UPA) inaccuracies, their advantages and disadvantages and their data requirements. Subsequently, we discuss the availability of three classes of data on which ACP optimization methods commonly rely in order to analyze the consequence of these data requirements. We then discuss the concepts of minimal perturbation and recommendation-based optimization and their addressing in existing literature. In addition, current shortcomings and research gaps are identified and avenues for future work can be highlighted. The discussion is presented in Section 5. Fig. 1 gives an overview of the general approach.

This work is focused on the Access Control Models (ACMs) Role-Based Access Control (RBAC) [18] and ABAC [19] since these are the most common ACMs. The standard RBAC model defines ACPs in the form of roles, which are bundles of permissions that can be assigned to subjects (e.g. employees) in a well-organized way. RBAC also offers the definition of constraints, which are statements that express negative authorizations, i.e. authorizations that must not be granted by the role set. ABAC ACPs in contrast are modeled as (dynamic) policies which make authorization decisions based on attributes of subjects (e.g. employees), objects (permissions) or the execution environment (e.g. the execution time). The term ACP hence refers to (constrained) RBAC roles as well as to ABAC policies [1]. The Organization for the Advancement of Structured Information Standards (OASIS) defined the eXtensible Access Control Markup Language (XACML) standard, which provides a notation for expressing ABAC policies in an XML-based format as well as a reference architecture for a policy evaluation mechanism [20]. XACML is the most important technical standard for ABAC and is explicitly addressed in many ABAC-related publications.

Note that ABAC was often proclaimed as the successor of RBAC² since ABAC has the descriptive strength to express RBAC along with other attribute-based policies. E.g. [21] showcase this behavior for an industrial use case. However, ABAC has still not reached the maturity of RBAC and has difficulties with practical adoption [22,23]. Most issues of ABAC are rooted in the raised flexibility of attributes and ABAC policies which backfire as increased complexity and thus are less comprehensive for administrators and policy engineers [22]. Additionally, spill-over effects from RBAC to ABAC and vice versa can be observed in literature [24–26]. Further well-known ACMs, like Discretionary Access Control (DAC) and Mandatory Access Control (MAC) [1], or very recent approaches like Attribute-aware Relationship-Based Access Control (AReBAC) [27–29] are not in the scope of this work.

The remainder of this paper is structured as follows. Section 2 presents work that is related to the quality of ACPs and their optimization. Section 3.2 defines ACP quality and ACP optimization and presents common quality criteria. Section 4 presents a literature survey on optimization of ACPs. Building on the findings of the survey, further aspects of ACP optimization are discussed in Section 5. Finally, Section 6 sums up the results and concludes this work.

2. Related work

While existing studies present a comprehensive picture of RBAC and ABAC research and open challenges [22,30], to the best of our knowledge no literature study has examined the more specific field of ACP optimization yet. Since publications that contribute to this field directly are presented and analyzed in the literature survey in Section 4, this chapter serves to present research areas that are closely related to ACP quality optimization.

The assessment of ACP quality is a fundamental building block for their optimization. Since ACP quality comprises many dimensions, the scientific literature is heterogeneous and proposes many ways to assess ACP quality. This includes the definition of quality metrics [7,31] and distinct research realms that aim at individual quality-related objectives. Examples are ACP anomaly analysis, which aims at the identification of conflicts and redundancies, or XACML evaluation runtime analysis, which aims to find causes for a slow evaluation runtime. Moreover, ACP quality research is concerned with the question of how the quality of ACPs develops in real environments and which structural reasons are responsible for this. Researchers have documented that

² In 2013 Gartner stated that “by 2020 70% of enterprises will use attribute based access control [...] as dominant mechanism to protect critical assets [...]” <https://www.gartner.com/en/documents/2607617>.

the quality of ACPs gradually deteriorates without targeted countermeasures, and have identified structural causes such as structurally determined over-granting [5] or the role explosion problem [32].

The initial creation of ACPs with high quality is addressed by ACP modeling approaches. Existing approaches differ greatly in terms of modeling objectives and their definitions of *optimal* ACPs [33,34]. Existing proposals typically focus on policy mining, i.e. the automatic generation of new ACPs based on *existing* authorization structures, or policy engineering, i.e. the (mostly manual) definition of new ACPs based on an *ideal* authorization structure. Hybrid approaches aim to utilize the advantages of mining and engineering [35,36]. While the initial creation of high-quality ACPs does not constitute an optimization, some RBAC or ABAC mining algorithms provide a maintenance mode that works on existing ACPs and were included in the literature survey. The repeated re-generation of an entire policy set is an alternative approach to keep ACPs up-to-date [37,38]. It is based on the assumptions that policies remain unchanged during their whole life cycle and can be generated on a sufficiently high quality level without any human interaction. Instead of maintaining existing ACPs, this approach aims to ensure a sufficient quality by re-generating all policies from scratch on a regular basis.

3. Theoretical foundations of ACP optimization

Before presenting the methodology and findings of the survey, this chapter introduces the theoretical foundations for ACP quality and its optimization. Section 3.1 demonstrates that ACP quality is an instance of the data quality concept as defined by Wand and Wang [39] and provides a definition of ACP quality based on it. To complement this definition, Section 3.2 presents a collection of 16 well-established ACP quality criteria. Building on these preliminaries, Section 3.3 concludes the theoretical foundation with a definition of ACP optimization.

3.1. Quality of access control policies

Present research on ACP quality shows similarities to data quality research: There is consensus in both fields that quality is a multidimensional concept and cannot simply be assessed as universally *good* or *bad*. As in the field of data quality research, IAM research also applies different criteria to assess the quality of ACPs. While some of them only make sense in the context of ACPs (like the grade of automation or evaluation runtime), others are equivalent to data quality dimensions that are well-established outside the research realm of IAM (e.g. accuracy, understandability, completeness or redundancy). To the best of our knowledge, present research did yet not formally show that ACP quality constitutes an instance of the data quality concept. In this section, we show that the widely acknowledged data quality model by Wand and Wang [39] is applicable to ACPs and provide a definition of ACP quality based on it.

The basis for an organization's access control decisions is its security policy [1,40]. A security policy is a collection of (often informal) requirements that define the authorizations of an organization. One of the most commonly quoted security policy requirements is the principle of least privilege, which states that a subject should not inherit more permissions than it needs to perform its tasks [10]. A security policy may also comprise requirements that are not directly related to the organization's security, for example to fulfill organizational or regulatory needs. One example are Segregation of Duty (SoD) requirements, which define mutually exclusive authorizations and are commonly employed to avoid conflicts of interest. Access Control Policies (ACPs) are machine-processable rules that can be evaluated in a fully automated manner to determine which accesses a subject is allowed to make [1]. Following these established definitions, ACPs are a data representation of the authorizations that are specified by an organization's security policy.

The data quality model by Wand and Wang [39] defines an information system as an entity which exists parallel to a real-world system. The data stored in the information system is a representation of a perception of the real-world system. Through interpretation of this data, a user perceives a view of the real-world system as inferred from the information system. The process of creating data that represents real-world entities is called the representation transformation. The process of creating an interpretation of the representational data which resembles the original real-world entities is called interpretation transformation. If both, representation transformation and interpretation transformation, work correctly, the view of the real-world system as inferred from the information system is identical to the view of the real-world system gained from direct observation. Any disparity in between these views represents a data deficiency. The authors define four assumptions which need to be met for the model to be applicable:

(i) *The Representation Assumption: An information system is a representation of a real-world system as perceived by users.* The authorizations defined by an organization's security policy exist outside the information system and are hence (abstract) real-world entities. The ACPs that express them are consequently a data representation of real-world entities, and the information system that stores them is a representation of a real-world system.

(ii) *The Interpretation Assumption: An information system is built for use by the user whose view of the real-world system is captured in the design of the system.* Wand and Wang explain that this assumption serves to ensure that the interpretation transformation (i.e. the process of transforming data back into perceivable real-world entities) will be able to map the data representation back to the original real-world entities. In the instance of ACPs, the design of the information system equals the view of the information system, because the representation transformation and the interpretation transformation are based on the same access control model: Any set of authorizations can be represented as a user-permission matrix. Both access control models in scope, RBAC and ABAC, allow to define an ACP set for every possible user-permission matrix, that will be mapped back the exact same user-permission matrix.

(iii) *The Inference Assumption: The information system can create a perceptible representation from which the user can infer a view of the real-world system as represented in the information system.* Since every ACP set can be represented (and also visualized) as a user-permission matrix, an information system can always infer a view of the original authorizations that are represented by the ACP set.

(iv) *The Internal View Assumption: Issues related to the external view such as why the data are needed and how they are used are not part of the model.* This assumption is self-fulfilling as it is merely states that the model does not deal with external issues.

By showing that ACPs fulfill these four assumptions, we show that the data quality model by Wand and Wang [39] is applicable to ACPs and that the concept of ACP quality constitutes an instance of the data quality concept. Present research widely agrees that the quality of data is best described as its "fitness for use" [41]. In accordance with this definition, we define the quality of ACPs as their fitness for use with regard to one or more quality dimensions that reflect the application context of access control.

3.2. Established quality criteria

Present research applies many different criteria to evaluate the quality of ACPs, many of which include a concept of optimal quality. Becklerle and Martucci [7] developed six criteria to determine well-usable ACPs sets and developed metrics to quantify these. Both Kunz et al. [33] and Mitra et al. [34] present surveys on role mining approaches and point out objectives that role mining algorithms apply to achieve high-quality roles. Jabal et al. [31] define a list of policy analysis criteria with implications to the policy quality. Besides that, a large number of both RBAC and ABAC related publications define quality criteria

“on the fly” and often also define metrics for these quality criteria to approach a particular objective at hand.

In order to complement the definition of ACP quality, the remainder of this section presents a collection of properties of ACPs that are commonly applied in existing literature to evaluate the quality of ACPs. The collection comprises 16 properties, 14 of which have an optimum in terms of ACP quality. The remaining two criteria (usage and relevance) are often used in the context of quality assessment, but are not an expression of ACP quality themselves. Furthermore, seven of the presented criteria affect the evaluation of ACPs by the access control mechanism directly, while the remaining nine only implicitly affect their correct evaluation through factors such as error-proneness during policy administration or maintenance efficiency. Please note that this is not a complete list, as creating a full list of established quality criteria would require a structured, reproducible literature survey on its own. Individual quality criteria may be positively or negatively correlated, or not influence each other at all: For example, adding new rules to an ABAC policy will likely increase its UPA coverage and bring it closer to being complete. At the same time, the complexity of the policy is increased, which suggests a negative correlation between the objectives of minimal complexity and maximal completeness. Kunz et al. [33] present a dependency analysis of quality criteria applied during role mining. Despite that, to the best of our knowledge, the interaction of ACP quality criteria has not been analyzed by scientific research yet. Table 1 provides an overview of all 16 properties, their optima and whether or not the criteria affect the ACP evaluation directly.

Accuracy is the most important quality dimension for ACPs as it expresses the effective correctness of their access control decisions. It is hence directly related to the correct evaluation of the policies. The accuracy of an ACP set defines, how accurately it represents the authorizations defined by the security policy³ [7]. There are two types of errors that decrease the accuracy of an ACP set: If an ACP set grants excessive UPAs, subjects inherit more permissions than they require. In contrast, missing UPAs mean that subjects require particular permissions, but are not granted them by the ACP set. The challenge in identifying inaccurate UPAs lies in determining which UPAs *should* be granted. If the entire set of correct UPAs was known, an ACP set could be optimized for perfect accuracy in a fully automated manner.

Excessive UPAs are a violation of the principle of least privilege and can cause a security vulnerability. An example of how harm can be done by excessive authorizations is when a hospital employee publishes sensitive patient data, either accidentally or in malicious intent. Excessive authorizations can also be abused by external attackers, for example, if an authorized employee is blackmailed or his or her user account is hijacked. The minimization of excessive UPAs is the primary objective when modeling and administering ACPs and is required to enable an acceptable level of security. Finding excessive UPAs however poses a greater challenge, and over-allocated privileges often go unnoticed until they are misused for a malicious act. For this reason, excessive UPAs tend to accumulate over time, making targeted countermeasures necessary [42]. The amount of excessive UPAs that are granted by an ACP set is hence a crucial indicator for its quality.

Missing UPAs keep subjects from doing their work and hence conflict with business continuity. For example, a company’s supply chain could suffer outages because an employee lacks the authorization to post a goods receipt in the enterprise resource planning system. Missing UPAs have a higher visibility than excessive UPAs because their damage occurs relatively quickly: A subject who has been wrongly deprived of an entitlement can immediately thereafter no longer perform a certain

Table 1
Common quality-related ACP properties.

Property	Optimum	Affects evaluation
Accuracy	Max	Yes
Excessive UPAs	Min	Yes
Missing UPAs	Min	Yes
Maintainability	Max	No
Understandability	Max	No
Sem. meaningfulness	Max	No
Complexity	Min	No
Redundancy	Min	No
Conflicts	Min	Yes
Grade of automation	Max	Yes
Evaluation runtime	Min	Yes
Similarity to opt. state	Max	No
Risk	Min	No
Completeness	Max	Yes
Usage	Ambiguous	No
Relevance	Ambiguous	No

task. The impact can be substantial since this can include very basic privileges, such as authorization to enter an organization’s premises or to log into their work station.

Maintainability describes how well an ACP set can be administered and kept up-to-date [43,44]. Low maintainability makes an ACP set prone to errors and leads to a higher administration effort. Increasing an ACP set’s maintainability is hence a prime objective of ACP optimization. The maintainability of an ACP set is influenced by several properties including its understandability, its complexity, its redundancy or the amount of conflicts that it contains, which can be assessed and optimized individually. Tool-supported administration of ACP also helps early on in keeping desired properties up-to-date [45].

Understandability expresses how well an ACP set can be understood by humans. It is closely related to maintainability and is often cited together [44,46]. A cryptic ACP set that is hard to understand is also hard to administer or maintain: For example, a policy administrator could misunderstand the meaning of an ABAC policy and make erroneous changes that cause inaccuracies. Alternatively, an administrator could decide not to make changes at all to a policy that he or she does not understand, leading to fast obsolescence of that policy. Maintaining a good understandability is hence considered a prime challenge by several authors [47,48]. A key factor for understandable ACPs is semantic meaningfulness. Moreover, several authors proposed approaches for visualizing ACPs, which aim to improve the understandability of an ACP set without applying any changes to it [23].

Semantic meaningfulness means that ACPs represent a human-understandable real-world concept. It is often argued to be crucial for ACP understandability [11,49]. Since semantic concepts can be described with attribute values, the semantic meaningfulness of an ACP can be assessed by measuring its accordance with semantically meaningful attributes [49]. For example, a role that can be described as “This role grants all permissions that are required for all software developers” would accord to 100% to the value “software developer” of the semantically meaningful employee attribute “job title” and hence has a very high semantic meaning. Note that this definition does not require the ACP to be defined based on attributes itself. Moreover, an attribute-based ACP does not automatically have a higher semantic meaning than a role, since it can define a long list of permitted or denied UPAs which share little or no semantic meaning.

Complexity expresses the amount of elements that an ACP set consists of. For example, an ABAC policy that comprises 200 statements is likely more complex than one that contains only 5 statements. Similarly, a set of 50 roles with hierarchies among them and numerous permission assignments is likely more complex than a set of 5 roles with few permission assignments and no role hierarchies. Low complexity improves the maintainability of an ACP set and reduces

³ Beckerle and Martucci [7] refer to the Security Policy as “Access Control Policy”. To avoid confusion, this work uses the more common term “Security Policy”. The term “Access Control Policy” refers to the data representation of the authorizations defined by the security policy, which is subject to optimization.

the computational effort required for its evaluation. Existing research assesses the complexity of an ACP set in numerous ways, including the amount of roles and ABAC policies contained in it, the size of roles and ABAC policies and more specific measurements. The most generic definition of ACP complexity is the Weighted Structural Complexity (WSC), which is a weighted sum of all elements defined by the underlying ACM [49,50]. The complexity of an ACP set is among the most commonly cited quality indicators. Unlike its understandability or semantic meaningfulness, the complexity of an ACP set can be quantified objectively without requiring any further data.

Redundancy occurs if an ACP set defines positive or negative authorizations more than once. For example, an employee could inherit the permission to close customer requests in a ticketing system twice because he has the role “customer support employee” and the role “administrator of ticketing system”. Similarly, a redundant negative authorization could occur for a bank employee if an ABAC policy defined that no bank employee who serves private customers may approve lending, and that no employee who is still in training may approve lending. Redundancy leads to an unnecessary bloating of ACP complexity. Moreover, ACP redundancy is a possible cause for administration errors, as a redundantly defined positive or negative authorization must be removed more than once for the change to become effective. If one of the redundant definitions is overlooked in the process, the ACP set obtains an inaccuracy. The impact can be substantial, for example when an emergency permission revocation process (i.e. a process where a subject is immediately stripped of all permissions, for example because the digital identity was stolen) fails because a redundant permission assignment is overlooked.

Conflicts exist within an ACP set if it defines both positive and negative authorizations for the same user-permission pairs. This means that a particular permission is both allowed and forbidden for the same user. For example, an ABAC policy could state that IT administrators have file access to an application server, while at the same time denying access to personnel data files for anyone outside the human resource department. While a static conflict is present within the effective UPAs that an ACP set realizes at a given time, dynamic conflicts are *potential* conflicts, i.e. conflicts that could arise due to the dynamic nature of the ACPs, but did not necessarily generate a contradiction yet [51]. While ABAC or XACML ACP sets define both positive and negative authorizations by default, RBAC ACP sets can only contain conflicts if used with RBAC constraints [18] since unconstrained roles define only positive authorizations. If a conflict exists within an ACP set, the access control mechanism must resolve it in order to make an unambiguous authorization decision. This is achieved by applying a conflict resolution strategy which defines how to make authorization decisions if a conflict occurs. The XACML standard defines basic conflict resolution strategies [52] and many more sophisticated conflict resolution algorithms were proposed by researchers. However, conflict resolution only aims to enable the access control mechanism to make a deterministic decision *despite* the presence of conflicts. It does not update the ACP set to remove the conflict and hence does not constitute a quality optimization. Since conflicts make an ACP set’s authorizations ambiguous, they are a possible cause for inaccuracies and reduce the ACP set’s understandability. Moreover, real-time conflict resolution reduces the evaluation performance of the ACP set.

The **grade of automation** of an ACP set determines, to which extent the authorizations defined by it adapt to new situations dynamically without requiring manual updates. For example, an ABAC policy that defines authorizations based on an employee’s department affiliation requires no updating of a policy definition if an employee moves into another department, since the employee’s department attribute value would change, thus leading to an updated result in the evaluation of subsequent authorization requests. Attribute-based ACPs inherently offer the possibility of dynamic rule definitions, since changes in referenced attribute values also change the authorization decisions resulting from ACPs evaluation. The standard RBAC model [18] in contrast

is static and cannot update authorizations automatically unless it is extended with an automation mechanism (such as [53–55]). While ACP automation is directly related to the policy evaluation, it is also a critical factor in ACP maintenance as it reduces administrative effort and prevents excessive and missing UPAs before they occur [13,42]. However, it cannot make ACP maintenance obsolete, since an automation mechanism operates with a limited scope, and dynamic ACP definitions can out-date or be erroneous like static ones (cmp. Section 4.3.6). Additionally, automation also eases other IAM processes like policy refinement, policy verification or conflict resolution [56,57].

The **evaluation runtime** of an ACP set determines how quickly it can be loaded and evaluated by an access control mechanism to answer an authorization request [58,59]. A sufficiently low evaluation runtime is critical if ACPs must be evaluated in real-time, since a pending authorization decision is a performance bottleneck for all relying application systems. As a result, users could spend considerable time waiting for simple button clicks to be executed, or performance-critical operations in an organization’s IT infrastructure such as large data processing tasks could pile up. The real-time evaluation of authorization requests by a central access control mechanism is a requirement specified in the OASIS XACML reference architecture [19]. The evaluation runtime of an ACP set can be influenced by many factors, like the amount of contained ACPs, or the amount, size, order of the rules that an ACP contains or the algorithm used [60,61].

Similarity expresses how similar one ACP set is compared to another. For example, a set of 10 department roles is likely more similar to a set of 5 department roles and 5 roles that represent an employee’s job title than it is to a set of 50 roles which represent employees’ job titles. Similarity is often used as a quality criterion by measuring how similar the assessed ACP set is to another ACP set which is considered optimal [62,63]. By maximizing its similarity to the optimal ACP set, a new ACP set can be generated that resembles the structure of an existing, well-structured ACP set, but includes the results of the newly applied generation algorithm. Moreover, similarity can be used to estimate the amount of updates that is required to migrate from one ACP set to another [64]. By applying high similarity to the original ACP set as an objective, an ACP optimization method can produce updates with *minimal perturbation*, which is expected to reduce the administrative effort required for implementing the changes.

Risk determines the impact of excessively assigned permissions for an ACP or an ACP set. For example, a policy that grants subjects full access to a banking system is more risky than a policy that grants subjects access to a WiFi hotspot. The risk of an ACP reflects the aggregated risk of the permissions that are permitted by it. ACP modeling or optimization methods can use risk minimization as an objective to reduce the impact of excessive permission assignments [65,66]. High risk is also an indicator for high maintenance priority and can serve as context information for policy engineers and reviewers, based on the assumption that a high risk value suggests a more restrictive handling than a low one [42].

Completeness determines the amount of UPAs that are covered by an ACP set. Examples for incomplete UPA sets could be a set of roles which contain only a subset of the permissions managed by an organization, or an ABAC policy that only contains attribute definitions to make authorization decisions for a subset of the requested subjects. If an ACP set does not cover UPAs, an access control mechanism is unable to determine an access control decision for the corresponding user-permission pair during policy evaluation (except for a standard fallback decision). The criterion of completeness is also often used in ACP modeling to determine to which extent a newly modeled ACP set expresses the UPAs defined by an input state [33]. The term coverage is closely related and often used interchangeably, but does not necessarily contain a quality indication. The coverage of an ACP is an indicator for its relevance, based on the assumption that a policy that defines many authorization decisions is more important than a policy that defines few ones.

Usage determines how often an ACP was invoked, i.e. how often the authorizations defined by an ACP were requested and executed. For example, an employee in the goods receiving department of a company might use the authorization to debit a delivery every day. An example of an infrequently used UPA could be that a back office employee needs to submit a balance sheet only once a year. If an UPA is never executed, this is an indicator that it is not needed, i.e. it is excessively assigned according to the principle of least privilege. The usage of an ACP can be reconstructed via access logs (as defined in Section 4.2.3), by counting the invocations of the UPAs that are covered by the ACP over a defined timespan. ACP usage can be used to determine UPA inaccuracies (cmp. Section 5.1.2) and is hence an important tool for ACP quality assessment. Moreover, the usage of an ACP is an indicator of its relevance, assuming that an ACP that is often invoked is more important than one that is scarcely invoked [16,67].

Relevance expresses how strongly an ACP influences authorization decisions in productive operations. For example, a basic role that allows every employee to access their work stations is likely more relevant than a specialist role that allows few employees to create a new email distribution list. ACP relevance is commonly analyzed to determine priorities when assessing or optimizing ACP quality. In case of conflicting policies, relevance can be used to prioritize a policy that should overrule another (cmp. Section 4.3.5). The relevance of a policy can also be an indicator of timeliness [68], and ACPs with low relevance can be interpreted as a security risk since they are likely to be over-permissive [31]. Moreover, policies with higher relevance can be maintained with higher priority in order to improve the effectiveness of ACP maintenance [69]. The properties *Risk*, *Completeness* and *Usage* are closely related to the relevance of an ACP and often used as an indicator.

3.3. Optimization of access control policies

Building on the definition of ACP quality, we define ACP optimization as an improvement of the quality of existing ACPs with regard to specified quality criteria. This definition has two important implications: First, since ACP quality is multidimensional, ACPs cannot be optimized towards a universal optimum, but only towards a particular optimization objective. For example, consider an update operation that assigns a new permission to a role in order to provide this permission to the employees which inherit the role. This optimization may reduce the amount of missing UPAs and hence constitutes an optimization with regards to accuracy. At the same time, adding a new assignment to the role set increases its complexity and hence reduces its quality with regards to this quality dimension. Second, ACP optimization requires that an existing ACP set is updated rather than a new ACP set being created from scratch. This implication seems obvious, as the quality of a data state can only be improved in comparison with a reference state, i.e. the original data state that existed before the optimization. However, since every update operation generates a new data state, every optimization of an existing ACP set could also be interpreted as the creation of a new ACP set. The optimization of an existing ACP set differs in its goal from modeling a new one in that it aims to leave the existing set structurally intact. This goal stems from the need to maintain ACPs at a high quality level over time at a reasonable cost: Daily operations, such as an employee's department change, or the integration of new application systems into an organization's IT infrastructure, require frequent updates to an ACP set. Both neglecting such updates and sub optimal updates have a negative impact on the quality of the ACP set. Without continued quality measures, it is common for ACP sets to proliferate over time, while changing real-world conditions and proprietary policy updates cause their quality to degrade [2–5]. Unlike ACP modeling, ACP optimization aims to make partial changes to a (possibly very large) ACP set while leaving the remainder of the ACP set unchanged. As a result, an ACP set that was modeled with high quality can be kept on high quality with significantly lower

effort than was required for its initial modeling. Moreover, by leaving an ACP set structurally intact, ACP optimization is able to retain the (often informal) semantic meaning of the existing ACPs that may be known only to human policy engineers. One possible approach to ensure that an ACP set remains structurally intact during optimization is the concept of “minimal perturbation”, which is discussed in detail in Section 5.3.

4. Literature survey on access control policy optimization

4.1. Survey methodology

The literature survey follows the methodology proposed by Levy and Ellis [17]. It aims to examine the existing body of research on ACP optimization (as defined in Section 3.3) and provide a structured analysis of the research field. Since ACP optimization has to follow defined optimization objectives, we worked out six optimization objectives that are relevant and well established in existing literature. These six optimization criteria are presented along with further categorization criteria in Section 4.2. We define the scope of literature included in the survey as “*Scientific publications that propose means for optimization of roles or attribute-based ACPs with regard to at least one of the six defined optimization objectives*”.

In accordance with the applied survey methodology, the literature research was started with a bibliographical database search. We used combinations of generic keywords like “maintenance”, “improvement”, “optimization”, “correction” and more specific keywords for the distinct optimization objectives, such as “redundant” or “redundancy”. Structured permutations of these keywords were entered in the online databases ACM Digital Library,⁴ IEEE Digital Library⁵ and Google Scholar.⁶ All publications that were not obviously related to another topic were added to an initial list of “unfiltered literature”. We also included all publications from the ACM Symposium on Access Control Models and Technologies (SACMAT)⁷ conference from the years 2001 to 2021 in this list. Every publication in the “unfiltered literature” list was then screened, which means that we read it superficially to determine whether it is relevant with regards to the survey topic. In this step, we read the title, abstract, introduction and conclusion and used the document search function to determine how the keywords applied during the bibliographical search were used. Doing so, we narrowed down the list of unfiltered literature to a second list of “relevant literature”, which needs to address either the quality of ACP in general or one of the optimization objectives. All publications that were regarded relevant for the topic were then read in depth to determine whether they fit the survey scope and could be added to the final list of “included literature”. This list formed the ACP optimization literature catalogue which was categorized and analyzed. Since a bibliographical search could only serve as an entry point, we conducted author and reference search for all publications in the list of “relevant literature” and added the resulting publications back into the initial list of “unfiltered literature”. Repeating the screening process for these publications, we executed a recursive search that allowed for a deep exploration of relevant research realms. A schematic overview of the applied literature research process is given in Fig. 2.

Due to its broad scope, the survey comprised a heterogeneous literature base. As a result, the literature was hard to grasp with keywords. The majority of relevant results was yielded via author and reference search. Many publications in the “included literature” list do not explicitly define ACP optimization as maintenance or quality optimization, but rather define distinct objectives with a narrow scope

⁴ <http://dl.acm.org/>.

⁵ <http://www.computer.org/>.

⁶ <http://scholar.google.com/>.

⁷ <http://www.sacmat.org/>.




Fig. 2. Applied literature research process in accordance with Levy and Ellis [17].

Table 2
Coverage of selected quality criteria in literature.

Quality criterion	Coverage
Excessive UPAs	Fuchs et al. [42], Hill [71], Jaferian et al. [14], Hummer et al. [46] and Puchta et al. [23]
Missing UPAs	Benedetti and Mori [43,72], Fuchs and Pernul [73], Colantonio et al. [74] and Meier et al. [48]
Redundancy	Guarnieri et al. [75], Shamooun et al. [76], Hu et al. [77], Mitra et al. [34] and Kunz et al. [33]
Conflicts	Hounder [78], Deng and Zhang [79], Dia and Farkas [80] and Shamooun et al. [76]
Complexity	Mitra et al. [34], Kunz et al. [33], Servos and Osborn [22], Currey et al. [81], Fuchs et al. [42] and Molloy et al. [11]
Grade of automation	Fuchs et al. [42], Kunz et al. [13], Hu et al. [82], Kern and Walhorn [53], Al-Kahtani and Sandhu [54] and Aftab et al. [55]

which are semantically equivalent. Moreover, many relevant research realms used other keywords than we would have expected: For example, the realm of “XACML anomaly analysis” aims at the identification (and sometimes removal) of anomalies in XACML policies, which can be either conflicts or redundancies and hence fit the survey scope. Since we did not know of this research realm beforehand, we hardly could have found it using keywords.

4.2. Research and categorization criteria

As seen in Section 3, ACP optimization is not universal, but can only improve the quality with respect to defined optimization objectives. To define the research scope of the survey, we screened existing literature and selected six optimization objectives that are central for ACPs’ fitness for use. These six optimization objectives serve as criteria for selecting relevant literature for this study (cmp. Section 4.1) and are presented in detail in the following Section 4.2.1.

Beside a textual analysis, the applied survey methodology suggests to categorize the literature catalogue. For a meaningful categorization, the selected criteria need to be concrete enough for in-depth insights while covering a heterogeneous literature catalogue. This leads us to three categories within this survey: (i) The first category is the targeted ACM of a publication. The majority of analyzed publications (except [46,70]) can be categorized as either RBAC or ABAC related. (ii) The next category is the optimization objective which serve simultaneously as research criteria. (iii) Another category is the contributed ACP optimization research artifact of the publication. (iv) Finally, the used data of the optimization method is analyzed as the last category. While the distinction between RBAC and ABAC is self-explanatory, the remaining categorization criteria are presented in Sections 4.2.1 to 4.2.3.

4.2.1. Optimization objective

Existing literature defines several quality criteria on the basis of which ACPs can be optimized (cmp. Section 3.2). To clearly define and narrow the scope of the survey, we searched the literature for optimization goals that are critical to the fitness for use of ACPs. Beckerle and Martucci [7] conduct semi-structured expert interviews and a literature analysis to identify critical requirements for obtaining usable ACP sets. Based on their results, they argue that the main aim of ACP optimization should be to improve the accuracy and maintainability of ACPs. They specify these requirements and work out six optimization objectives that serve these two goals. The authors also develop metrics for the quantification of these criteria and conduct two user studies to evaluate them. To the best of our knowledge, this is the only scientific publication that documents a structured research process

for developing ACP optimization objectives and provides a conclusive evaluation. The optimization criteria are: (i) “Allow no more than the owner wants to be allowed”, (ii) “Allow everything the owner wants to be allowed”, (iii) “A rule must not be fully covered by another rule of the same rule set”, (iv) “Two rules belonging to the same rule set must not conflict”, (v) “Minimize the number of rule set elements”. and (vi) “Minimize maintenance effort in a changing system”. Each of these six optimization criteria can be mapped to one ACP quality dimension that was presented in Section 3.2. We reformulate them to accord with the common literature terminology based on the addressed quality dimensions: (i) *Reduce excessive UPAs*, (ii) *Reduce missing UPAs*, (iii) *Reduce redundancy*, (iv) *Reduce conflicts*, (v) *Reduce complexity* and (vi) *Increase grade of automation*. We conducted a literature search to confirm the relevance of these six quality dimensions. Table 2 lists further publications that underline their importance for ACPs’ fitness for use. Note that this list is not exhaustive. Due to their thorough foundation in existing literature and their frequently argued importance for optimization, we selected these six optimization objectives as the basis for the literature survey.

4.2.2. Research artifact

The analyzed publications on ACP optimization are heterogeneous not only in terms of the addressed optimization objectives, but also in terms of their contributions. We identified four types of research artifacts that are repeatedly presented to contribute to the optimization of ACPs: (i) Optimization process models analyze ACP optimization from a business perspective. They define process steps, roles and responsibilities and analyze how the technical optimization can be embedded into a changing real-world environment. (ii) Optimization algorithms define formal ways to modify ACPs in order to improve their quality. They receive an existing ACP set (and possibly supplemental data) as input and generate an optimized ACP set as output. (iii) Optimization tools aid humans in the semi-automated optimization of ACPs. Publications with this kind of contribution present optimization tools to demonstrate how tool-supported ACP optimization can be done. (iv) ACM extensions propose ways to enhance ACMs to improve ACP quality. In the analyzed literature catalogue, this type of research artifact is only present in the form of RBAC extensions that enhance role definitions in order to provide automation. Note that we were not able to categorize the research artifact of a publication into one of these standardized categories in 5 cases.

4.2.3. Data usage

For the last criteria group, the survey literature was analyzed for the kind of data that was used to perform ACP optimization. We applied the conceptual IAM data model proposed by Kunz et al. [47] to achieve an

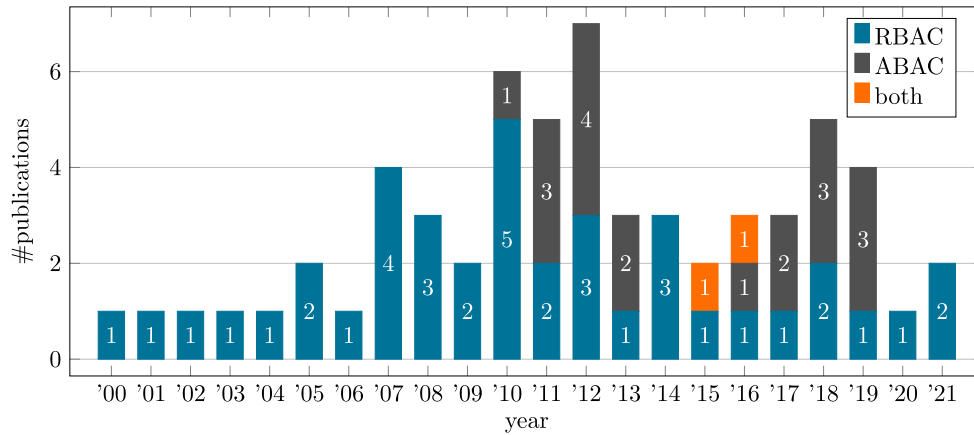


Fig. 3. Publications over time.

integrated view on the processed data. The model defines the central data entities that are processed in IAM and their relations towards one another. In addition to RBAC and ABAC, it integrates the conceptual entities defined in six central IAM technology standards⁸ and proposes a terminology that is suited to cover the integrated concepts. According to this definition, a digital identity is a representation of a human user, which can be a record of an employee that is stored and processed in a human resource management system. Within an application system, digital identities possess accounts, to which permissions are assigned. Permissions can be hierarchically nested, in which case they are inherited transitively. Parallel to the definitions within an application system, digital identities can be assigned roles (as defined by RBAC), which inherit permissions through assignment and can also be hierarchically nested. Beside that, permissions can be granted via policies, which are ABAC or XACML policies in the context of this work. Note that in our terminology, both roles and policies are ACPs. The model defines a context entity, which is a scenario that can be evaluated by a policy (for example environmental conditions in an ABAC policy). At last, the model defines an attribute entity, which expresses a property of a digital identity, an account or a permission and can be evaluated by policies.

We identified a total of five types of data that were repeatedly used for ACP optimization: The *ACP set*, a *user-permission matrix*, *entity attributes*, *access logs* and *update logs*. From this set, we chose the use of entity attributes, access logs and update logs as categorization criteria. Processing of the ACP set and the user-permission matrix was not analyzed since these data types are trivial: The existing ACP set is the most basic type of data and must be known in order to be optimized. It comprises the roles or ABAC policies that are in effect and are updated in the course of the optimization. A user-permission matrix (also called access control matrix or UPA set) is the most basic representation of the permission assignments that an ACP set grants [11]. It is a boolean matrix which holds a true or false value for every possible user-permission combination. A user-permission matrix is commonly visualized in the form of an access grid and is not limited to a particular ACM [48]. An optimization that updates an ACP set needs to verify that it did not create inaccuracies. Since it is the most basic permission assignment information available, a user-permission matrix is commonly assumed to be available for any ACP optimization effort.

⁸ Lightweight Directory Access Protocol (LDAP), Security Assertion Markup Language (SAML)/Shibboleth, Service Provisioning Markup Language (SPML), Open Authorization (OAuth), System for Cross-domain Identity Management (SCIM) and XACML.

Entity attributes are properties of entities other than the optimized ACPs themselves. As defined by Kunz et al. commonly expressed entities are digital identities, accounts and permissions. Attributes are often used to give ACPs semantic meaning: Since an attribute reflects a real-world property, binding an ACP to a given attribute can bind it to its semantic meaning [11]. Optimization approaches that require entity attributes need to process actual value expressions of concrete entities. In contrast, a method that restructures a given XACML policy by reordering its attribute statements without processing any entity attribute values would not be classified into this category. The reliance on entity attributes is a limiting factor since their availability may be limited. Moreover, the use of entity attributes for ACP optimization means that the results are dependent on the value of said attributes at the time of the optimization, meaning that an optimization result may lose validity when the attribute values change.

Access logs express historic accesses of users to permissions. While their notations differ, access logs can be displayed as a tuple $\langle S, O, A, R \rangle$ that represents a historic access request, with S being the requesting subject, O being the requested object, A being the requested action and R the result of the request, i.e. *permit* or *deny* [5]. Note that some publications only consider successful permission invocations, thus reducing access logs to a tuple $\langle S, O, A \rangle$ of permitted access requests. Access logs provide valuable insight on the actual need of permissions and can help to identify missing or excessive permission assignments.

Update logs are the second type of historic data used for ACP optimization. Update logs contain information on past changes of IAM related entities, for example a modification of an ACP, the creation of a new user account or the change of an employee's department affiliation. Update logs can be used to identify real-world events that provide important ACP update information (for example, the job change of an employee might require a change of his or her permissions) and can provide insight on the development of an ACP set over time. To the best of our knowledge, no scientific publication exists that defines the structure of IAM update logs.

4.3. Criteria-based analysis

4.3.1. Overview

The literature survey yielded 61 publications that provide means for optimization of existing RBAC and ABAC ACPs since the year 2000. Out of these publications, 42 address the optimization of roles, 21 address the optimization of attribute-based policies, and two address both ACMS. The optimization of ACPs has been addressed continuously over the past 20 years. While the peak of interest for RBAC optimization occurred between 2007 and 2014, the topic receives steady attention to this day, underlining that role optimization remains a relevant research

Table 3
Categorized literature for ABAC optimization [3,44,46,62,63,67,69,70,75–80,83–89].

Publication	Optimization objective					Research artifact				Data usage			
	Reduce excessive UPAs	Reduce missing UPAs	Reduce complexity	Reduce redundancy	Reduce conflicts	Increase grade of automation	Optimization process model	Optimization algorithm	Optimization tool	ACM extension	Entity attributes	Access logs	Update logs
Argento et al. [83]	✓							✓			✓	✓	
Benkaouz et al. [84]			✓					✓					
Cheng et al. [44]					✓								
Deng and Zhang [79]					✓			✓					
Dia and Farkas [80]					✓			✓					
Guarnieri et al. [75]				✓				✓					
Hadj et al. [63]			✓	✓				✓					
Hadj et al. [67]				✓	✓			✓				✓	
Hadj et al. [85]				✓	✓			✓					
Hadj et al. [69]				✓	✓			✓				✓	
Hein et al. [86]					✓			✓					✓
Hounder [78]					✓			✓					
Hu et al. [3]				✓									
Hu et al. [87]				✓	✓			✓					
Hu et al. [77]				✓	✓			✓					
Hummer et al. [46]	✓	✓					✓				✓	✓	
Hummer et al. [70]	✓	✓					✓				✓	✓	
Narouei and Takabi [62]			✓					✓			✓		
Oberholzer [88]			✓					✓					
Shamoon et al. [76]					✓			✓					
Stepien et al. [89]			✓	✓				✓					

subject. The first analyzed ABAC optimization paper was published in 2010. The Tables 3 and 4 present all analyzed publications and their categorization according to the survey criteria defined in Section 4.2. Fig. 3 depicts all analyzed sources ordered by their year of publication.

4.3.2. Considered optimization scenarios

To connect ACP optimization with real-world scenarios, it is helpful to consider optimization scenarios. While various themes for optimization scenarios are possible, e.g. IAM goals [118] or usage of popular techniques like access reviews [15], it is reasonable to take a closer look at maturity and automation as its driver. Like shown by Schimpf et al. [119] a driver for higher maturity is automation which requires or builds upon underlying optimization methods. In the sense of automation, ACP optimization scenarios can thus be considered as manual, semi-automated and automated.

In manual optimization, a human administrator or policy engineer updates part of an ACP set based in their individual context knowledge. Manual adjustments to an ACP set are commonly done in daily operations, e.g. to grant employees new permissions after their responsibilities changed. Another example are access reviews, where a responsible human (e.g. a department head) tries to find excessively assigned permissions manually and marks them for revocation [14]. Due to the lack of automation, manual optimization is limited to small ACP sets or subsets of larger ACP sets. Another optimization scenario is the fully automated updating of an ACP set. In this scenario, an ACP optimization algorithm generates optimization steps for an ACP set. The resulting changes will be implemented in the underlying applications automatically. This approach has the disadvantage that an algorithm has no understanding of the semantic structure of an ACP set. As a result, the structure of the optimized ACPs may differ greatly from their original structure. In addition,

fully automated methods can hardly be integrated into established change management processes, which require that changes to ACPs must be confirmed by a responsible employee. The third scenario is semi-automated optimization. Semi-automated processes try to bridge the gap between manual and automatic optimization by enabling humans with technical support to optimize a very large set of ACPs. A common example are recommendation-based optimization methods, which generate possible ACP updates automatically, and delegate them to responsible humans for decision. Updates will only become effective if the responsible human agrees to them. Semi-automated optimization can be supplemented by ACP visualization procedures and other data analysis techniques which are not optimization methods themselves. The concept of recommendation-based optimization is discussed in Section 5.4.

The publications analyzed during the survey cannot always be clearly assigned to one of these scenarios. We observe that publications that propose an optimization algorithm typically assume fully automated optimization. Still, some optimization algorithms allow that the resulting change steps are delegated to humans in the form of recommendations, thus enabling them for semi-automated optimization [43,72,82,96,106]. Publications that present an optimization tool aim to support humans in the semi-automated optimization of ACPs by definition. Of the publications analyzed that present a process model, all except [111] assume that changes to an ACP set can be generated by algorithms, but must be delegated to a human for decision before becoming effective. It is noticeable that publications that aim to embed ACP optimization into an organization's processes largely reject fully automated optimization (except [111]). Publications that propose an extension of an ACM define ways to extend the structure of a role set with automation logic. However, they do not define a concrete optimization scenario in which a role set is updated to inherit this automation logic. ACM extensions are thus rather a blueprint for ACP optimization and not limited to a concrete optimization scenario.

Table 4
Categorized literature for RBAC optimization [2,11,12,15,16,42,43,46,53,55,64,70,72,82,90–117].

Publication	Optimization objective						Research artifact				Data usage		
	Reduce excessive UPAs	Reduce missing UPAs	Reduce complexity	Reduce redundancy	Reduce conflicts	Increase grade of automation	Optimization process model	Optimization algorithm	Optimization tool	ACM extension	Entity attributes	Access logs	Update logs
Aftab et al. [55]						✓				✓	✓		
Al-Kahtani and Sandhu [90]						✓				✓	✓		
Al-Kahtani and Sandhu [91]						✓				✓	✓		
Baumgrass [92]	✓	✓										✓	
Benedetti and Mori [72]		✓	✓				✓	✓				✓	
Benedetti and Mori [43]	✓	✓	✓				✓	✓				✓	
Chakraborty and Ray [93]						✓				✓	✓		
Fuchs et al. [42]	✓	✓	✓				✓					✓	✓
Gal-Oz et al. [94]				✓				✓					
Groll et al. [15]	✓										✓		✓
Guo et al. [95]			✓					✓					
Han et al. [96]						✓		✓		✓	✓		
Herzberg et al. [97]						✓				✓	✓		
Hu et al. [82]						✓	✓		✓				
Hu et al. [98]						✓		✓					
Hu et al. [99]						✓		✓					
Huang et al. [100]						✓		✓		✓	✓		
Hummer et al. [46]	✓	✓					✓				✓	✓	
Hummer et al. [70]	✓	✓					✓				✓	✓	
Kern and Walhorn [53]						✓				✓	✓		
Lu et al. [101]						✓				✓	✓		
Lu et al. [102]						✓				✓	✓		
Molloy et al. [11]			✓	✓				✓			✓		
Ni et al. [103]						✓				✓	✓		
Pan et al. [16]			✓					✓				✓	
Pang et al. [104]				✓	✓			✓					
Pang et al. [105]			✓	✓				✓					
Parkinson et al. [12]						✓				✓	✓		
Rao et al. [106]						✓		✓		✓	✓	✓	
Saffarian et al. [107]						✓				✓	✓		
Shafiq et al. [108]					✓			✓					
Sheng and Osborn [109]						✓				✓	✓		
Strembeck [110]				✓					✓				
Strembeck [111]				✓			✓					✓	✓
Takabi et al. [112]						✓		✓		✓	✓		
Takabi and Joshi [113]			✓					✓					
Vaidya et al. [64]						✓		✓					
Xia et al. [2]			✓					✓					
Yi-qun et al. [114]						✓				✓	✓		
Zhang et al. [115]			✓					✓					
Zhang et al. [116]	✓	✓						✓				✓	
Zong et al. [117]						✓				✓	✓		

4.3.3. Reduce excessive & missing UPAs

We identified nine publications that aim at identifying and rectifying inaccurate UPAs. Out of these nine publications, six address both missing and excessive permission assignments. Another two exclusively aim at excessive permission assignments and one addresses only missing permission assignments. Fuchs et al. [42] propose a process model for RBAC optimization. Hummer et al. [46,70] present a process model for both RBAC and ABAC optimization. Benedetti and Mori [72] present both an RBAC optimization process model and a

Max-SAT algorithm that uses access logs to identify missing permission assignments and adjust roles while minimizing their complexity. They expand it in a follow-up publication to address excessive permissions as well [43]. Baumgrass [92] and Zhang et al. [116] both use access logs to identify missing or excessive UPAs and adjust roles accordingly. Argento et al. [83] use access logs to identify excessive permission assignments and update ABAC policies. Groll et al. [15] propose to use negative access review decisions, i.e. decisions in which a human reviewer identified excessive UPAs, to identify similar UPAs and generate revocation recommendations for them, thus amplifying

the impact of manual identification of excessive permissions. To the best of our knowledge, this is the only publication proposing a method that uses a data source other than access logs to identify excessive permissions automatically. Note that approaches which provide automation for RBAC updating are also relevant for this optimization objective as they aid the closely related objective of *preventing* UPA inaccuracies. These approaches are analyzed in Section 4.3.6.

The key challenge for correcting missing or excessive UPAs is to identify them. Once found, such inaccuracies can be corrected in a fully automated manner. Throughout the literature analysis we identified three basic approaches for identifying missing or excessive permission assignments: First, manual identification requires that a human overlooks an ACP set and tries to find assignments which he or she knows to be inaccurate. This task is commonly executed in the form of access reviews, which are the de-facto standard process for this task and mandated by central regulation frameworks. Second, usage-based approaches analyze access logs for information on historic permission invocations and can thus be used to determine the set of actually needed UPAs. Third, update history based approaches analyze historic updates of the ACP set to identify missing or excessive permission assignments. The advantages and data requirements of these basic concepts are discussed in Section 5.1.

4.3.4. Reduce complexity

The survey yielded ten publications that aim to reduce the complexity of an RBAC state and five publications that aim to simplify an ABAC state. Out of the RBAC related publications, Fuchs et al. [42] are the only one that addresses the problem exclusively at the process level. They propose a process model for the optimization of roles regarding accuracy and complexity. Benedetti and Mori [43,72] define an RBAC maintenance process model and an algorithm that optimizes the complexity of an RBAC state. Xia et al. [2] define the Role Refinement Problem and provide an algorithmic analysis, but also define process-level requirements for role refinement. The remaining six RBAC-related publications address complexity optimization strictly on a technical or algorithmic perspective. Takabi and Joshi [113], Pan et al. [16] and Benedetti and Mori [43,72] aim to minimize the amount of updates required to perform the generated optimizations. Molloy et al. [11] use user attributes to increase the semantic meaning of the roles that are either mined or maintained in their approach. They also argue for the use of access logs and update logs, but do not use them themselves. Pan et al. [16] use access logs to calculate the usage of roles, which they use to determine its relevance as one indicator of its quality. The five ABAC-related publications propose algorithms to reduce the complexity [62,84,88] or both complexity and redundancy [63,89] of an existing ABAC state. Hadj et al. [63] and Narouei and Takabi [62] calculate the similarity of input and output state to minimize the required updates. Narouei and Takabi [62] require user and permission attributes for the refinement of an ABAC policy set. Altogether, ACP optimization for complexity requires no data other than the ACP set and the user-permission matrix since it is a mere algorithmic problem that can be solved without semantic knowledge. As long as the resulting UPAs remain unchanged, the ACP set can be rearranged for a lower complexity in a fully automated manner. However, context knowledge (for example provided by access logs) can provide further insight into the meaning of ACP and benefit the optimization efforts.

4.3.5. Reduce redundancy

We identified 14 publications that propose methods for reducing redundancy in existing roles or attribute-based policies. Strembeck [110] define a high-level maintenance process for roles that aims at reducing redundancy. Strembeck [111] present a role engineering tool that supports humans in finding and removing redundant role-permission assignments. The remaining twelve publications propose algorithms for the reduction of redundancy in roles or XACML policies. We observe

that optimization methodologies often consider redundancy reduction a secondary optimization, as four of the analyzed publications aim to reduce complexity and five rectify conflicts as their primary objective. While the reduction of redundancy is a non-trivial algorithmic problem, it can be addressed very well with algorithms since solving it at core requires no data other than the optimized ACP set and its UPA set. Consequently, most of the analyzed approaches use no further data. However, Molloy et al. [11] use entity attributes to mine new and optimize existing roles with redundancy as one optimization objective. Strembeck [110,111] propose the use of context constraints during role optimization and define a process for engineering it, using access logs and update logs as data sources.

4.3.6. Reduce conflicts

A large research area aims at the resolution of ACP conflicts. However, most of these publications propose methods to generate a permit or deny decision during the evaluation of ACPs despite the existence of a conflict. They do not modify the underlying ACPs to correct the error and hence do not fit into the scope of this survey. Altogether, we identified 14 publications that propose methods for the reduction of ACP conflicts. Cheng et al. [44] define a methodology for removing inconsistencies (which include conflicts) in rule-based ACPs such as ABAC policies. Despite their method being named “removing process”, we chose not to categorize it as a process model as it is a technical methodology that is closer to defining an algorithm than a (business) process model. The remaining 13 publications define conflict reduction algorithms. Both Pang et al. [104] and Shafiq et al. [108] propose graph optimization algorithms to remove constraint conflicts on an RBAC state. The remaining eleven publications propose algorithms for the reduction of rule conflicts in an XACML policy set.

A crucial challenge with direct impact to the ACP quality is to decide whether a given conflict should be converted into a permit or deny decision before updating the ACP set accordingly. A wrong correction decision would still remove a conflict in the ACP set, but lead to a UPA inaccuracy in return. The analyzed literature offers different approaches to make this decision. Hu et al. [3,87] and Deng and Zhang [79] embed existing conflict resolution strategies in their approach. Hu et al. [77] refine the previously proposed approach to apply conflict resolution strategies for individual segments of an ACP set, and apply removability constraints to individual ACPs in order to achieve a more fine-grained correction decision. Both Dia and Farkas [80] and Hadj et al. [85] extend ACPs with scope constraints to support the correction decision. Dia and Farkas [80] also generate recommendations for decisions on conflict removal which can be delegated to responsible entity owners if the correction algorithm generated a non-mandatory update. Hein et al. [86] generate update logs that track the changes to an ACP and use these in order to support administrators and allow them to restore a previous state which is known to be correct. Shafiq et al. [108] apply priorities to role constraints and evaluate these during the conflict correction. At last, Hounder [78] defines algorithms to aid human policy administrators in conflict correction, thus proposing a semi-automation of conflict correction which can incorporate the semantic knowledge of a human in the resolution decision.

We observed that conflict correction research primarily focuses on attribute-based policies with only two publications for role conflict correction yielded by the survey. We explain this by the fact that attribute-based ACPs are prone to conflicts since they define both positive and negative authorizations by default. Roles themselves in contrast define only positive authorizations and can only be conflicted if used with constraints [120]. Nevertheless, role constraints are crucial part of role-based access control and are required to express role restrictions for fundamental process-level requirements such SoD policies. The current state of research on RBAC conflict correction does not reflect this.

4.3.7. Increase grade of automation

Since ABAC is dynamic by nature, providing automation is only relevant for RBAC in the scope of the survey. Out of 22 publications that provide means for RBAC automation, we group 16 publications into three classes: Rule-based automation, learning-based automation and trust-based automation.

Rule-based automation extends roles with rules which evaluate attributes to automatically update employee-role assignments [53,55,90,96,100,114], role-permission assignments [12,55,96,100] or role hierarchy assignments [91]. They can hence be interpreted as hybrids of RBAC and ABAC. Since attributes can reflect a real-world property with semantic meaning, attribute-based role automation can help to improve the semantic meaningfulness of roles. A single rule can cover a multitude of role model updates and can hence greatly reduce administrative effort for a role set. Nevertheless, automation rules can outdate over time and need to be maintained together with the role set.

Learning-based automation approaches use existing entitlement information as input to find valid role updates. Sheng and Osborn [109] use entity attributes to generate employee-role updates and Ni et al. [103] use entity attributes to generate role-permission updates. Rao et al. [106] utilize access logs to automatize employee role assignment. Learning-based automation works without static rules and hence does not necessarily require manual definitions. Unlike rule-based automation, learning-based automation may react to changes that the role engineers did not consider when modeling the role set. However, learning-based automation requires learning data as input, and the quality of its updates is limited by the quality of the input data. As a result, we argue that learning-based optimization might be better suited to maintain a role set with an already high quality, than to optimize a role set with low quality from scratch.

Trust-based automation assigns roles to users based on the users' trustworthiness [93,97,107,112,117]. Trust-based approaches differ from the two previously presented ones as they do not assign roles based on the tasks that a user has to perform (as required by the principle of least privilege), but try to assess a set of maximum permissible permissions. This is typically done by calculating trust scores for users and defining minimum trust scores that a user needs to have in order to be granted certain roles or permissions. Trust-based automation approaches are designed for open environments where the users are not completely known (for example collaborative platforms like Wikis). They are not designed for classical inhouse identity management environments where employees with a defined task range are managed (cmp. [121,122]), but aim to mitigate risk when little user information is available.

We identified another six publications that provide means for automatic role updating, but do not try to identify possible updates themselves: Hu et al. [98] generate migration paths to automatize the implementation of role model updates in application systems. Hu et al. [82] define a tool and a process which aid the automatic updating of role-permission assignments by checking whether an update is achievable with a given set of constraints. Lu et al. [101] and Hu et al. [99] evaluate role updating algorithmically to determine the complexity of automatic checking for role-permission and role-role assignments. Lu et al. [102] propose a role generalization algorithm that aims to optimize roles for automatic assignment via user authentication queries. Vaidya et al. [64] aim to enable role mining algorithms for optimization of an existing role set by generating a state that is as similar as possible to an existing role set and an optimal one.

Altogether, we conclude that RBAC automation is thoroughly covered by research. It can help to reduce administrative effort and maintain a high role model quality. However, it does not make role maintenance obsolete, since all types of automation have limitations and cannot be expected to react to all future changes adequately. Beside a limited scope, RBAC automation configurations themselves can be erroneous or outdate over time just like a classical role set.

5. Discussion

Building on the results of the literature survey, this chapter discusses important aspects of ACP optimization. In doing so, we analyze several concepts that are repeatedly addressed in the survey literature and play a critical role in the optimization of ACPs. Since the identification of UPA inaccuracies is the most critical challenge in correcting excessive and missing UPAs, we work out three prototypical approaches commonly found in the literature and discuss their advantages and limitations as well as their data requirements in Section 5.1. We then examine the availability of the three types of data that were included in the literature analysis in Section 5.2. Subsequently we discuss the concepts of minimal perturbation and of recommendation-based optimization in Sections 5.3 and 5.4. At last the limitations of this work are discussed in Section 5.5.

5.1. Identification of UPA inaccuracies

Finding inaccurate UPAs is the key challenge when optimizing ACPs for accuracy. If the complete set of required UPAs is known, an ACP set can be optimized for perfect accuracy in a fully automated manner. However, such a set does not usually exist in an explicit form, and finding out which permissions any subject should be granted in accordance with an organization's security policy is difficult and time-consuming. ACP accuracy optimization methods hence face the primary challenge to identify as many inaccuracies as possible in order to be able to correct them. During the literature analysis we identified three prototypical approaches to identify UPA inaccuracies, which we present and discuss below.

5.1.1. Manual identification

Manual identification primarily concerns excessive UPAs, because manual identification of missing UPAs can simply be initiated by the affected users, for example by ordering a missing permission in a structured process through an IAM entitlement shop [123]. The most obvious approach to identify excessive UPAs is to have a human who knows an organization's security policy check the effective UPAs and search for inaccuracies. A process that works by this scheme is known by the name of *Access Reviews*. Access Reviews are a standard IAM process that is executed periodically and aims to identify and rectify excessive UPAs (and sometimes other data inaccuracies like inaccurate attribute values) [14]. Their execution is strongly driven by external requirements raised by compliance frameworks or IT security standards [121,124]. During an Access Review, a responsible human reviews a list of entitlement assignments for the users, ACPs or permissions within their responsibility, and decides whether they are still necessary. Based on this decision, the reviewer will either confirm an assignment's correctness, or refuse to confirm it, in which case the assigned permissions will be revoked. Access Reviews are a central measure to prevent the accumulation of excessive permissions [14]. However, they are error-prone and tend to overlook and confirm excessive UPAs due to a number of structural challenges [15]. In particular, the sheer amount of data that has to be processed during Access Reviews can be overwhelming, and the decision whether a user requires a particular permission is difficult to make. In case of uncertainty, Access Review decisions are biased towards confirming an existing assignment for a number of reasons: (i) Since an existing UPA likely has been subject to an approval process before becoming effective, a reviewer has reason to believe that an existing assignment has a legitimate reason. (ii) If a reviewer makes a false decision, only a false revocation of an existing UPA would have an immediate consequence, because an employee would no longer be able to execute a certain task as a result. In contrast, an erroneous confirmation of an already granted UPA is unlikely to have an immediate effect as long as the resulting security vulnerability is not abused for malicious action. (iii) For the same reason, the visibility of an erroneous revocation within an

organization is higher than the visibility of an erroneous confirmation, which adds a social incentive for reviewers to simply confirm existing UPAs as to avoid visible errors. As a result, the effectivity of Access Reviews is limited. While some research effort tries to aid users in the effective execution of Access Reviews [14,125], to the best of our knowledge only one approach was proposed that aims to measure the quality of Access Review decisions in order to identify erroneous UPA confirmations automatically [15]. Despite the importance of Access Reviews and the difficulties practitioners face in implementing them, the tasks of aiding companies and users in the execution of Access Reviews and measuring their effectivity were scarcely addressed by research.

5.1.2. Usage based identification

Usage based identification of UPA inaccuracies evaluates historic permission invocations to determine which permissions a particular user actually needed in the past. It is hence directly related to the principle of least privilege, which states that any user should not inherit more permission than necessary to perform his or her tasks. Several ACP optimization approaches use access logs for usage based identification of UPA inaccuracies (cmp. Section 4.3.3). If access logs are available, excessively assigned permissions can be found relatively easy by comparing historic permission usages with the actually granted permissions of an ACP set. Any permissions that are granted for a user, but were not used within a specified time frame (e.g. over the last year) can be interpreted as excessively assigned and are hence candidates for removal. This approach is well suited to detect large quantities of excessively assigned permissions that were overlooked during UPA maintenance. However, it can struggle to distinguish excessive UPAs from accurate, but rarely used ones (e.g. a yearly creation of a report). Moreover, this approach can only identify excessive UPAs as long as they are not used (possibly even with malicious intent).

If access logs do not only include historic permission invocations, but also access requests which have been denied, they can be used to identify missing UPAs: If a permission invocation is often requested and denied, this can be an indicator that the permission is handled too restrictively. However, the identification of missing UPAs leaves more room for interpretation than the identification of excessive UPAs and requires finding a balance between business continuity considerations and security considerations. We hence argue that the paramount value of access logs lies in identifying excessive UPAs. Especially UPA accumulation, i.e. the accumulation of permission assignments that were once granted but are no longer required by a user (for example because the user's responsibilities have changed since the permission assignment) is a major problem that can be addressed very well with the described approach. However, usage based correction is limited by the availability of access logs, which is discussed in Section 5.2.

5.1.3. Update history based identification

The use of update histories for the optimization of ACP accuracy was proposed by several authors. Some also provided examples how update histories could be used for this cause. Fuchs et al. [42] cite the update history of a role model as an important source of context information for optimizing a role model. However, they only consider it for manual evaluation by human role engineers. Strembeck [111] names trace management as necessary requirement for maintaining complex models, but does not go into the details of ACP maintenance in this regard. Hein et al. [86] propose an algorithm to generate update logs that track the changes made to an ACP set and use them to enable administrators to perform a rollback operation and restore previous states. Mitra et al. [34] propose to use update logs to identify roles which decayed over time in order to select candidates for maintenance. While the authors do not elaborate on this, using historical data to assess timeliness is common practice in other areas of data quality research [126]. Molloy et al. [11] name the event pattern of a user

losing several permissions, followed by being assigned several permissions within a short time frame as an indicator for a job change event and argue that such events provide valuable context information for the creation of high-quality roles. While not explicitly naming it as a role optimization use case, they also stress that the creation of high-quality roles is not a once-and-for-all effort, but must be succeeded by continuous role optimization efforts. They provide further examples for meaningful events within a role model's update history, arguing that permissions which are often assigned or unassigned together are likely related to the same real-world context, thus providing an example of using update logs to determine the semantic meaning of permissions. Molloy et al. [11] also argue that historic update information provides evidence on legacy permissions which should be removed from the role model. To the best of our knowledge, no methods were proposed that use update logs for the correction of missing or excessive UPAs. However, Groll et al. [15] propose an approach to find erroneous Access Review decisions (i.e. decisions where an excessive permission assignment was falsely confirmed) automatically. The proposed approach, which uses revocation decisions from the analyzed Access Review as input, performs learning-based outlier detection to find other Access Review decisions which are likely to be over-permissive. Although the method does not explicitly use historic update information, it can be abstracted as an analysis of UPA unassignment events, which could also be taken from update logs instead of Access Review decisions. Altogether we conclude that ACP update histories are a promising source of context information for identifying excessive or missing UPAs. To this day, existing research has made few attempts to make use of this information source.

5.2. Data availability

To highlight the limitations resulting from the data requirements of the analyzed ACP optimization approaches, we will discuss the availability of the investigated data types. As defined in Section 4.2 we assume the conceptual IAM model of Kunz et al. [47] for an integrated view on the processed data. The three classes of data for which we analyzed the survey literature catalogue are entity attributes, access logs and update logs. First of all, the use of all of these data sources requires an integrated view on the processed IAM data. This requirement is not trivial since an identity management infrastructure comprises a wide range of heterogeneous data sources. These data sources may either provide a centralized data view (e.g. a human resource system that stores the employee data for the entire organization or a directory system that serves as a centralized user account and permission data storage) or exist as numerous decentralized data sources that have different data schemes and data storages (e.g. application systems that manage individual user accounts and permissions for their own application context) [121]. While having an integrated data view on the identity management infrastructure is a common prerequisite for IAM measures, the creation of such a view requires significant effort and should therefore not be neglected [122].

5.2.1. Availability of attributes

Attributes are properties of entities within the IAM data view other than the ACPs themselves. If an integrated view of the related entities exists, entity attributes are from a technical point of view easily obtainable. Due to the sensitivity of personal data however, the processing of attributes of digital identities may be restricted. While this means that sensitive attributes (like an employee's loan details or sick leave history) might not be available for ACP optimization, it should not prevent attributes of digital identities from being processed altogether, as many central business-related attributes (such as departmental affiliation and job title) do not fall into this category.

5.2.2. Availability of access logs

The availability of access logs is more complex. The OASIS XACML reference architecture presumes that authorization requests are sent to a central access control mechanism with a Policy Decision Point (PDP) at its core. If this is the case, access requests (both granted and denied ones) can be logged completely. However, we argue that this is a prototypical architecture which comes with several challenges: (i) Creating a central access control mechanism requires that all applications which are subject to IAM efforts delegate their authorization decisions to a PDP. This requires significant integration effort, which may not be economically feasible. (ii) The central evaluation of access requests requires, that every authorizable action of a user within an application (which may in case of highly configurable systems come down to every single click on a button) be sent to the PDP, evaluated and answered in real-time before the user action is executed. Since this has to happen during the run time of the application with negligible time delay, this requires a highly traffic resistant, performant and available IT infrastructure and almost immediate access request evaluation on the PDP side, again implying efforts that may very well exceed the benefits of centralizing authorization decisions. (iii) This approach requires that all applications which are subject to IAM are technically able to delegate every authorization decision to a PDP, which is currently not the standard.

The alternative approach to obtain access logs is by decentralized logging within the applications where the permission invocations occur. The decentralized logs must then be collected and integrated in a central log stream. While the decentralized approach does not require real-time decision delegation, it also requires high integration effort and is limited by the ability of all managed applications to log permission invocations: Although several industry software solutions do provide access logs or access statistics,⁹ this is not a primary use case for many software products and hence not the standard.

Beside technical and economical limitations, access logs contain particularly sensitive personal data as they enable work monitoring, and their use may be restricted due to requirements of the legislator or employee representatives. This problem is also known in other fields which monitor IT infrastructure for security, for example in the context of Security Information and Event Management (SIEM) [127]. Due to the sensitivity of the processed data, it may be necessary to limit surveillance to the most critical areas, for example, users with a particularly high number of privileges (like administrators) or critical applications (like a banking system). Overall, we conclude that access logs are difficult to develop as a data source. In addition to significant technical hurdles, the economic viability of monitoring activities across the board is not always given. Moreover, ethical and legal hurdles must be considered. For these reasons, we argue that the availability of access logs that include all user accounts and permissions from all the applications in the scope of an organization's IAM cannot be assumed to be standard. Approaches that rely on the availability of access logs for optimization of ACPs may have limited applicability in practice.

5.2.3. Availability of update logs

Update Logs contain the update history of the entities processed in IAM. All availability limitations that apply to entity attributes hence also apply to entity attributes within the update logs. Apart from that, update logs are easy to obtain: If an integrated view of the IAM data exists, then its changes can be monitored, too. The creation of update logs is a common functionality for industrial IAM systems, which have a central view over IAM entities since they are used to manage ACPs and to provision them to the related application systems. Moreover,

⁹ For example, Microsoft Exchange (<https://www.microsoft.com/de-de/microsoft-365/exchange/email>) provides detailed logs on email distribution list usages and SAP ERP (<https://www.sap.com/>) provides aggregated statistics of transaction invocations.

legal regulations imply the (partial) existence of update logs: In order to check compliance with the principle of least privilege, an auditor must be able to investigate when and how an ACP set was updated to grant users new permissions. For this reason, we argue that update logs are a readily accessible data source. Nevertheless, to the best of our knowledge, no scientific model exists that describes how IAM update logs are structured. Foundation work is still missing for the scientific development of this data source and possible applications.

5.3. Minimal perturbation

The concept of minimal perturbation aims to optimize ACPs with as few changes as possible [64]. ACP optimization algorithms often address this goal by defining maximal similarity to the original state as a secondary optimization objective. The ability to optimize ACPs with few changes is an important factor that can determine the practicality of an optimization method for two reasons: (i) The fewer changes needed to achieve quality improvement, the lower the administrative effort required to implement those changes. Since many organizations are mandated to execute approval processes for changes to the ACP set, an ACP update often requires the interaction of humans (e.g. a role owner or employee owner) before it can be enforced. Moreover, since access control is often enforced decentralized within the application systems managed in an identity management infrastructure, changes to the ACP set have to be provisioned into the related application systems in order to take effect. Minimizing the perturbation hence improves the economic viability of ACP optimization. (ii) The fewer changes needed to achieve quality improvement, the higher the degree to which an ACP set remains structurally intact. ACPs are typically modeled to reflect semantic concepts. This goes from single policies reflecting simple statements (like "any employee of this organization may access the WiFi hot-spot") up the whole ACP set, which can be designed, for example, to reflect the organizational structure of a company [36,49]. Furthermore, ACPs incorporate (often informal) contextual knowledge that may only be known to human policy engineers. By largely preserving the structure of an ACP set after its initial creation, the semantic meaning of the ACPs and the contextual knowledge that has gone into them are also preserved, which is beneficial for the quality of the resulting ACP set and prevents it from becoming unrecognizable after repeated execution of optimization methods. Since minimal perturbation can be critical for the practical viability of an optimization method, we argue that it is not adequately addressed by existing research: Out of 38 publications in the literature survey catalogue that propose an algorithm for ACP optimization, only 10 consider minimal perturbation when defining their optimization objectives [16,43,62–64,72,82,106,113,116].

5.4. Recommendation-based optimization

In this section, we discuss the concept of recommendation-based optimization. Recommendation-based optimization means that ACP update steps produced by an optimization method are not applied directly to the underlying ACP set, but bundled into optimization recommendations. An optimization recommendation represents one or more changes to the ACP set and can be delegated for decision to responsible human decision makers. Only when the decision maker agrees to the optimization, the associated update steps are applied to the ACPs set. If no manual decision is required, an optimization recommendation can alternatively also be executed automatically.

The use of recommendations addresses business-related constraints that can be crucial for the practical applicability of an optimization method: (i) Approval processes often require that changes to the policy set be approved by a responsible subject (like a department head, an application owner or policy owner) before being implemented. By bundling optimization decisions into human-decidable steps, organizations can use (semi-)automatized optimization methods while

remaining compliant. (ii) Business constraints may exist (like regulatory requirements or practical hurdles) that prohibit some changes in the policy set, but are unknown to optimization algorithms. Moreover, automatic optimization methods struggle to understand the semantic structure and context meaning of an ACP set. As a result, updates may be proposed that are technically valid, but make little sense in the real world. With recommendations, a human decider can serve as a quality gate for optimization steps and prevent changes that are problematic or forbidden. Since domain experts are likely to have specific context knowledge, including them into the optimization process can improve the effectiveness of the optimization altogether. We argue that recommendation-based ACP optimization constitutes a hybrid model of fully-automated and manual optimization. In the closely related domains of RBAC and ABAC policy modeling, which face the closely related challenge to automate the creation of semantically meaningful ACPs with high quality, hybrid approaches have also been proposed and gained broad acceptance [36,128].

A number of optimization methods analyzed in the survey rely on recommendations. Fuchs et al. [42] propose a process model for RBAC optimization that includes a mechanism for generating new role extensions (i.e. new assignments of roles to employees, permissions or other roles). They define that every role extension is delegated to a role owner for decision. Hummer et al. [46,70] propose a process model for the optimization of RBAC or ABAC ACPs based on usage patterns. Similar to Fuchs et al. they define that every optimization step is processed via a recommendation mechanism. Benedetti and Mori [72] define an RBAC maintenance process that identifies missing role-permission assignments algorithmically. In [43], this process is extended to also find excessive permission assignments. They stress that the found “violations” must be presented to a security administrator who may confirm or reject them before they are processed to generate RBAC model updates. This approach differs from the previous recommendation-based optimization approaches as the recommendation mechanism is active in a preliminary stage of the optimization process, not at the end of it. Groll et al. [15] define a methodology that analyzes confirmation decisions of access reviews to find possible errors. The results are then delegated to a human policy analyst for inspection. If the analyst confirms the error, the underlying UPA is revoked, which means that every found possible error is equivalent to a recommendation to revoke the underlying UPA. Baumgrass [92] propose a process-centric methodology for refining existing RBAC states. They use event logs to derive RBAC artifacts (i.e. components of a role model like employee-role, role-role or role-permission assignments) to extend a role set. The authors stress that these artifacts are merely update candidates and must be integrated into the role model manually by a human decider, e.g. with the help of a role engineering tool. Hu et al. [82] present a tool which generates update steps that migrate an ACP set into a target state. Their tool can recommend different migration paths to a human, who has to decide which (if any) of them might be suitable. Rao et al. [106], Han et al. [96] and Chakraborty and Ray [93] propose RBAC extensions which generate recommendations for updating user-role assignments. While Rao et al. [106] and Han et al. [96] propose rule-based approaches, Chakraborty and Ray [93] propose to recommend user-role-assignments on the basis of users’ trustworthiness.

It is noticeable that especially those publications that take a process perspective in optimizing ACPs require that optimizations are recommended to human deciders before they are implemented. Publications that propose algorithms for optimization often leave this requirement out. As a result, many optimization methods are not applicable for recommendation-based optimization. In order to be used for recommendation-based optimization, an optimization method must fulfill three requirements: (i) An optimization method needs to create individually decidable ACP update steps. This means that it must be possible to discard one step without having to discard the remaining ones. (ii) The decidable update steps must be small enough to be

meaningfully decidable by a single decider. (iii) The decidable update steps must be human-understandable. The requirements (ii) and (iii) indicate that small update steps are preferable to large ones. Furthermore, it is helpful if the optimization method has a concept of the semantic structure of the ACP set. For example, ACPs that affect specific application systems can be treated and recommended in a bundle that is delegated to the respective system owner for decision. These requirements show that not every optimization algorithm can be reasonably adapted for recommendation-based optimization: For example, the output of a graph optimization algorithm that takes a role set as input and generates a single optimized role set cannot easily be converted into recommendations, since the decider would have to accept or reject the entire optimized role set. Otherwise, if only one included update step were rejected, the remaining optimization would be incomplete and the resulting role set would likely be erroneous. In a real-world business environment with requirements for business continuity, change management processes and regulatory requirements, the ability to embed ACP optimization steps into the existing process landscape is a crucial factory for the practical applicability of an optimization method.

5.5. Limitations of this work

The content of this study is limited by the selection of ACP quality and optimization criteria. The 16 quality criteria presented in Section 3.2 provide a broad overview of the quality consideration of ACPs in the scientific IAM literature. Nevertheless, we do not claim completeness, as a representative summary would require a structured, reproducible literature survey. Moreover, the research scope of the literature survey is limited by the six selected optimization objectives defined in Section 4.2.1. To ensure the relevance of the chosen criteria, we adopted the six optimization criteria developed by Beckerle and Martucci [7] and confirmed them as central to ACPs’ fitness for use in further literature research. Nevertheless, this represents a preselection that influences the literature unearthed in the survey and the findings based on it.

The literature research process of the survey also represents a possible limitation: We found that many publications use proprietary terminology, or aim at specific problems that are semantically equivalent to ACP maintenance or optimization, but not formulated as such (for example “ACP anomaly resolution”). The heterogeneity of the literature poses a challenge to structured literature research. Although we have adhered to the methodology presented in Section 4.1, we cannot rule out with certainty that we have overlooked literature that fits the scope of the survey.

Due to its broad scope, this work cannot provide a detailed comparative evaluation of the various approaches that researchers proposed for optimization of ACPs. The analyzed publications differ in terms of their contributed research artifact, their grade of automation, and their optimization objective. It is difficult to compare a process model for ACP optimization with an algorithm for ACP optimization, or to compare an algorithm for reducing redundancies with an algorithm for correcting conflicts. We tried to address this challenge by formulating a generalized problem description, and by formulating analysis criteria are applicable among the heterogeneous literature base. Moreover, this work does not evaluate the analyzed approaches for their effectiveness or correctness, since a formal evaluation of the analyzed publications would have exceeded its scope.

Furthermore, the quality of ACPs is not well developed scientifically. While existing research proposes numerous proprietary definitions of *good* ACPs, there are few publications that address the topic holistically. The six optimization objectives on which the survey is based are well documented in the literature. Other quality-related ACP properties that are presented in Section 3.2 are cited less frequently. Further groundwork on ACP quality is desirable.

6. Conclusion

This work studied the optimization of ACPs with the following contributions: (i) We show that the quality of ACPs constitutes an instance of the data quality concept as defined by Wand and Wang [39] and provide a definition of ACP quality based on it. We give a broad overview of ACP quality as perceived in the IAM literature and provide a definition of ACP optimization. (ii) We present a structured literature survey that categorizes and analyzes existing methods for ACP optimization. We point out that the reduction of excessive and missing UPAs comes with the paramount challenge of *identifying* these inaccuracies. Once found, such inaccuracies can be corrected in a fully automated manner. Since identifying UPA inaccuracies relies heavily on context data, the availability of such data is a bottleneck. The reduction of complexity and redundancy in contrast are algorithmic problems which can be solved without any data other than the existing ACP set and its user-permission matrix. Approaches for conflict correction face the challenge of deciding whether to resolve a conflict by granting or permitting a particular UPA, and existing literature offers many strategies to address it. The issue of providing automation for role updating is thoroughly covered. Many approaches can be categorized into one of three classes: Rule-based automation, learning-based automation and trust-based automation. (iii) Building on the structured survey, we analyze important aspects of ACP optimization in more detail. In particular, we discuss three basic concepts for identifying UPA inaccuracies and their advantages and limitations. We analyze three prototypical types of data on which ACP optimization methods commonly rely and discuss their availability: While many ACP optimization methods require the existence of access logs, we point out that their availability may be limited in practice. Update logs, on the other hand, have not yet been precisely defined by researchers or incorporated into optimization methods. However, many authors emphasize their value as a possible source of information in optimizing ACPs. Furthermore, we analyze the concepts of minimal perturbation and recommendation-based optimization and argue their relevance for ACP optimization. Although both concepts are known to the research community, we point out that most optimization methods do not take them into account.

Future work has several starting points to contribute to the optimization of ACPs. (i) The quality of ACPs is often addressed in research, but it lacks theoretical foundations. There is no comprehensive literature survey that compiles existing ACP quality dimensions in a structured and reproducible process (with the exception of the quality criteria works named in Section 3.2). Also, the correlations of quality dimensions, i.e. which dimension influences another positively, negatively or not at all, are not comprehensively studied. Furthermore, there is no model that analyzes which ACP quality dimensions are subordinate or superordinate to others: For example, accuracy or maintainability seem to be aggregated dimensions that are based on different properties of ACPs. A structured overview, e.g. in the form of a topology, could provide clarity here and help standardize terminology in IAM research. (ii) Update logs are hardly developed as a data source. While especially the identification of UPA inaccuracies urgently needs other data sources than access logs, there is no definition of how update logs are structured or can be obtained in an IAM infrastructure. This groundwork could be used to develop methods for the (semi-)automated identification of UPA accuracies, which could be used to improve security and business continuity in organizations. (iii) While purely algorithmic or purely process-related optimization methods have been proposed by research, there is a lack of publications describing how algorithmic optimization can be embedded in existing process landscapes. Methodologies that implement optimization procedures, e.g. with tool support, could help to close this gap. Furthermore, practical reports or case studies would be a valuable aid for researchers trying to align theoretical ACP optimization methods with real-world needs. An analysis of the extent to which fully or semi-automated optimization

is applicable in practical scenarios would also be helpful. (iv) Finally, there is very little research on access reviews: Although numerous organizations need to invest significant effort into the execution of access reviews on a regular basis, little research effort has been made to improve their limited effectiveness.

The optimization of ACPs is a relevant and ongoing research topic, and practitioners are left with major challenges that are yet to be solved. Existing research differs greatly in scope and terminology, which indicates that the research topic has yet to gain maturity. We hope to contribute to its standardization with this work.

CRedit authorship contribution statement

Sascha Kern: Conceptualization, Methodology, Validation, Data curation, Writing – original draft, Writing – review & editing, Project administration. **Thomas Baumer:** Data curation, Writing – original draft, Writing – review & editing, Visualization. **Sebastian Groll:** Conceptualization, Writing – review & editing. **Ludwig Fuchs:** Supervision. **Günther Pernul:** Writing – review & editing, Supervision, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The research leading to these results was supported by the German Federal Ministry of Education and Research as part of the DEVISE project (<https://devise.ur.de/>).

References

- [1] Samarati P, de Vimercati SC. Access control: Policies, models, and mechanisms. In: Focardi R, Gorrieri R, editors. Foundations of security analysis and design. Berlin, Heidelberg: Springer Berlin Heidelberg; 2001, p. 137–96.
- [2] Xia H, Dawande M, Mookerjee V. Role refinement in access control: Model and analysis. INFORMS J Comput 2014;26(4):866–84. <http://dx.doi.org/10.1287/ijoc.2014.0603>.
- [3] Hu H, Ahn G-J, Kulkarni K. Anomaly discovery and resolution in web access control policies. In: Proceedings of the 16th ACM symposium on access control models and technologies. SACMAT '11, New York, NY, USA: Association for Computing Machinery; 2011, p. 165–74. <http://dx.doi.org/10.1145/1998441.1998472>.
- [4] Xu T, Naing HM, Lu L, Zhou Y. How do system administrators resolve access-denied issues in the real world? In: Proceedings of the 2017 CHI conference on human factors in computing systems. New York, NY, USA: Association for Computing Machinery; 2017, p. 348–61, URL: <https://doi.org/10.1145/3025453.3025999>.
- [5] Xiang C, Wu Y, Shen B, Shen M, Huang H, Xu T, Zhou Y, Moore C, Jin X, Sheng T. Towards continuous access control validation and forensics. In: Proceedings of the 2019 ACM SIGSAC conference on computer and communications security. CCS '19, New York, NY, USA: Association for Computing Machinery; 2019, p. 113–29. <http://dx.doi.org/10.1145/3319535.3363191>.
- [6] Bauer L, Cranor LF, Reeder RW, Reiter MK, Vaniea K. Real life challenges in access-control management. In: Proceedings of the SIGCHI conference on human factors in computing systems. CHI '09, New York, NY, USA: Association for Computing Machinery; 2009, p. 899–908. <http://dx.doi.org/10.1145/1518701.1518838>.
- [7] Beckerle M, Martucci LA. Formal definitions for usable access control rule sets from goals to metrics. In: Proceedings of the ninth symposium on usable privacy and security. 2013, p. 1–11.
- [8] Tsiostas D, Kittes G, Chouliaras N, Kantavelou I, Maglaras L, Douligeris C, Vlachos V. The insider threat: Reasons, effects and mitigation techniques. In: 24th Pan-Hellenic conference on informatics. PCI 2020, New York, NY, USA: Association for Computing Machinery; 2020, p. 340–5. <http://dx.doi.org/10.1145/3437120.3437336>.
- [9] Gilbert N. 31 Crucial insider threat statistics: 2021 latest trends & challenges. 2021, URL: <https://financesonline.com/insider-threat-statistics/>.

- [10] Horne D. Permissions. In: Encyclopedia of cryptography and security. Boston, MA: Springer US; 2011, p. 924–7. http://dx.doi.org/10.1007/978-1-4419-5906-5_786.
- [11] Molloy I, Chen H, Li T, Wang Q, Li N, Bertino E, Calo S, Lobo J. Mining roles with multiple objectives. *ACM Trans Inf Syst Secur* 2010;13(4):1–35. <http://dx.doi.org/10.1145/1880022.1880030>.
- [12] Parkinson S, Khan S, Chrapa L. Automated planning for administrating role-based access control. AAAI; 2020.
- [13] Kunz M, Fuchs L, Hummer M, Pernul G. Introducing dynamic identity and access management in organizations. In: Information systems security. Springer International Publishing; 2015, p. 139–58. http://dx.doi.org/10.1007/978-3-319-26961-0_9.
- [14] Jaferian P, Rashtian H, Beznosov K. To authorize or not authorize: Helping users review access policies in organizations. In: Proceedings of the Tenth USENIX conference on usable privacy and security. SOUPS '14, USA: USENIX Association; 2014, p. 301–20.
- [15] Groll S, Kern S, Fuchs L, Pernul G. Monitoring access reviews by crowd labelling. In: Trust, privacy and security in digital business. Springer International Publishing; 2021, p. 3–17. http://dx.doi.org/10.1007/978-3-030-86586-3_1.
- [16] Pan N, Sun L, He L-S, Zhu Z-Q. An approach for hierarchical RBAC re-configuration with minimal perturbation. *IEEE Access* 2018;6:40389–99. <http://dx.doi.org/10.1109/access.2017.2782838>.
- [17] Levy Y, Ellis TJ. A systems approach to conduct an effective literature review in support of information systems research. *Inf Sci Int J Emerg Transdiscipl* 2006;9:181–212.
- [18] Sandhu RS. Role-based access control. In: Advances in computers. Elsevier; 1998, p. 237–86. [http://dx.doi.org/10.1016/s0065-2458\(08\)60206-5](http://dx.doi.org/10.1016/s0065-2458(08)60206-5).
- [19] Hu VC, Ferraiolo D, Kuhn R, Schnitzer A, Sandlin K, Miller R, Scarfone K. Guide to attribute based access control (ABAC) definition and considerations. Technical Report, U.S. Department of Commerce; 2014. <http://dx.doi.org/10.6028/nist.sp.800-162>.
- [20] Godik S, Moses T. 2003, URL: <https://www.oasis-open.org/committees/xacml/repository/cs-xacml-specification-1.1.pdf>.
- [21] Cheminod M, Durante L, Valenza F, Valenzano A. Toward attribute-based access control policy in industrial networked systems. In: 2018 14th IEEE international workshop on factory communication systems (WFCS). 2018, p. 1–9. <http://dx.doi.org/10.1109/WFCS.2018.8402339>.
- [22] Servos D, Osborn SL. Current research and open problems in attribute-based access control. *ACM Comput Surv* 2017;49(4). <http://dx.doi.org/10.1145/3007204>.
- [23] Puchta A, Böhm F, Pernul G. Contributing to current challenges in identity and access management with visual analytics. In: Data and applications security and privacy XXXIII. Springer International Publishing; 2019, p. 221–39. http://dx.doi.org/10.1007/978-3-030-22479-0_12.
- [24] Gardiyawasa Pussewalage HS, Oleshchuk VA. Attribute based access control scheme with controlled access delegation for collaborative E-health environments. *J Inf Secur Appl* 2017;37:50–64. <http://dx.doi.org/10.1016/j.jisa.2017.10.004>, URL: <https://www.sciencedirect.com/science/article/pii/S221421261730128X>.
- [25] Qi H, Di X, Li J. Formal definition and analysis of access control model based on role and attribute. *J Inf Secur Appl* 2018;43:53–60. <http://dx.doi.org/10.1016/j.jisa.2018.09.001>, URL: <https://www.sciencedirect.com/science/article/pii/S221421261730368X>.
- [26] Nazerian F, Motameni H, Nematzadeh H. Emergency role-based access control (E-RBAC) and analysis of model specifications with alloy. *J Inf Secur Appl* 2019;45:131–42. <http://dx.doi.org/10.1016/j.jisa.2019.01.008>, URL: <https://www.sciencedirect.com/science/article/pii/S2214212618303843>.
- [27] Cheng Y, Park J, Sandhu R. Attribute-aware relationship-based access control for online social networks. In: Atluri V, Pernul G, editors. Data and applications security and privacy XXVIII. Berlin, Heidelberg: Springer Berlin Heidelberg; 2014, p. 292–306.
- [28] Chakraborty S, Sandhu R. On feasibility of attribute-aware relationship-based access control policy mining. In: Barker K, Ghazinour K, editors. Data and applications security and privacy XXXV. Cham: Springer International Publishing; 2021, p. 393–405.
- [29] Chakraborty S, Sandhu R. Formal analysis of ReBAC policy mining feasibility. In: Proceedings of the eleventh ACM conference on data and application security and privacy. CODASPY '21, New York, NY, USA: Association for Computing Machinery; 2021, p. 197–207. <http://dx.doi.org/10.1145/3422337.3447828>.
- [30] Fuchs L, Pernul G, Sandhu R. Roles in information security – A survey and classification of the research area. *Comput Secur* 2011;30(8):748–69. <http://dx.doi.org/10.1016/j.cose.2011.08.002>, URL: <https://www.sciencedirect.com/science/article/pii/S016740481100099X>.
- [31] Jabal AA, Davari M, Bertino E, Makaya C, Calo S, Verma D, Russo A, Williams C. Methods and tools for policy analysis. *ACM Comput Surv* 2019;51(6):1–35. <http://dx.doi.org/10.1145/3295749>.
- [32] Elliott A, Knight S. Role explosion: Acknowledging the problem.. In: Arabia HR, Reza H, Deligiannidis L, Cuadrado-Gallego JJ, Schmidt V, Solo AMG, editors. Software engineering research and practice. CSREA Press; 2010, p. 349–55, URL: <http://dblp.uni-trier.de/db/conf/serp/serp2010.html#ElliottKI0>.
- [33] Kunz M, Fuchs L, Netter M, Pernul G. How to discover high-quality roles? A survey and dependency analysis of quality criteria in role mining. In: Communications in computer and information science. Springer International Publishing; 2015, p. 49–67. http://dx.doi.org/10.1007/978-3-319-27668-7_4.
- [34] Mitra B, Sural S, Vaidya J, Atluri V. A survey of role mining. *ACM Comput Surv* 2016;48(4):1–37. <http://dx.doi.org/10.1145/2871148>.
- [35] Cotrini C, Weghorn T, Basin D. Mining ABAC rules from sparse logs. In: 2018 IEEE European symposium on security and privacy (EuroSP). 2018, p. 31–46. <http://dx.doi.org/10.1109/EuroSP.2018.00011>.
- [36] Fuchs L, Pernul G. HyDro – hybrid development of roles. In: Information systems security. Springer Berlin Heidelberg; 2008, p. 287–302. http://dx.doi.org/10.1007/978-3-540-89862-7_24.
- [37] Calo S, Manotas I, de Mel G, Cunningham D, Law M, Verma D, Russo A, Bertino E. AGENP: An ASGrammar-based GENerative policy framework. In: Policy-based autonomic data governance. Springer International Publishing; 2019, p. 3–20. http://dx.doi.org/10.1007/978-3-030-17277-0_1.
- [38] Verma D, Calo S, Witherspoon SA, Manotas I, Bertino E, Jabal AA, de Mel GR, Swami A, Cirincione G, Pearson G. Managing training data from untrusted partners using self-generating policies. In: Pham T, editor. Artificial intelligence and machine learning for multi-domain operations applications. SPIE; 2019, p. 1–15. <http://dx.doi.org/10.1117/12.2519682>.
- [39] Wand Y, Wang RY. Anchoring data quality dimensions in ontological foundations. *Commun ACM* 1996;39(11):86–95.
- [40] Sandhu RS, Samarati P. Access control: principle and practice. *IEEE Commun Mag* 1994;32(9):40–8.
- [41] Tayi GK, Ballou DP. Examining data quality. *Commun ACM* 1998;41(2):54–7.
- [42] Fuchs L, Kunz M, Pernul G. Role model optimization for secure role-based identity management. In: European conference on information systems (ECIS). 2014, p. 1–15.
- [43] Benedetti M, Mori M. On the use of max-SAT and PDDL in RBAC maintenance. *Cybersecurity* 2019;2(1). <http://dx.doi.org/10.1186/s42400-019-0036-9>.
- [44] Cheng Z, Royer J-C, Tisi M. Removing problems in rule-based policies. In: ICT systems security and privacy protection. Springer International Publishing; 2019, p. 120–33. http://dx.doi.org/10.1007/978-3-030-22312-0_9.
- [45] Seifermann S, Heinrich R, Werle D, Reussner R. Detecting violations of access control and information flow policies in data flow diagrams. *J Syst Softw* 2022;184(C). <http://dx.doi.org/10.1016/j.jss.2021.111138>.
- [46] Hummer M, Kunz M, Netter M, Fuchs L, Pernul G. Advanced identity and access policy management using contextual data. In: 2015 10th international conference on availability, reliability and security. 2015, p. 40–9. <http://dx.doi.org/10.1109/ARES.2015.40>.
- [47] Kunz M, Puchta A, Groll S, Fuchs L, Pernul G. Attribute quality management for dynamic identity and access management. *J Inf Secur Appl* 2019;44:64–79.
- [48] Meier S, Fuchs L, Pernul G. Managing the access grid - a process view to minimize insider misuse risks. In: 11th international conference on wirtschafsinformatik (WI2013). 2013, p. 1051–65, URL: <https://epub.uni-regensburg.de/27930/>.
- [49] Xu Z. Mining meaningful role-based and attribute-based access control policies (Ph.D. thesis), Stony Brook University; 2014.
- [50] Molloy I, Chen H, Li T, Wang Q, Li N, Bertino E, Calo S, Lobo J. Mining roles with semantic meanings. In: Proceedings of the 13th ACM symposium on access control models and technologies. SACMAT '08, New York, NY, USA: Association for Computing Machinery; 2008, p. 21–30. <http://dx.doi.org/10.1145/1377836.1377840>.
- [51] Dunlop N, Indulska J, Raymond K. Methods for conflict resolution in policy-based management systems. In: Seventh IEEE international enterprise distributed object computing conference, 2003. Proceedings. IEEE; 2003, p. 98–109.
- [52] Moses T. Extensible access control markup language (XACML) version 2.0. 2005, URL: http://docs.oasis-open.org/xacml/2.0/access_control-xacml-2.0-core-spec-os.pdf.
- [53] Kern A, Walhorn C. Rule support for role-based access control. In: Proceedings of the tenth ACM symposium on access control models and technologies. SACMAT '05, New York, NY, USA: Association for Computing Machinery; 2005, p. 130–8. <http://dx.doi.org/10.1145/1063979.1064002>.
- [54] Al-Kahtani M, Sandhu R. Rule-based RBAC with negative authorization. In: 20th annual computer security applications conference. 2004, p. 405–15. <http://dx.doi.org/10.1109/CSAC.2004.32>.
- [55] Aftab MU, Habib MA, Mehmood N, Aslam M, Irfan M. Attributed role based access control model. In: 2015 conference on information assurance and cyber security (CIACS). 2015, p. 83–9. <http://dx.doi.org/10.1109/CIACS.2015.7395571>.
- [56] Cheminod M, Durante L, Seno L, Valenza F, Valenzano A. Automated fixing of access policy implementation in industrial networked systems. In: 2017 IEEE 13th international workshop on factory communication systems (WFCS). 2017, p. 1–9. <http://dx.doi.org/10.1109/WFCS.2017.7991947>.
- [57] Cheminod M, Durante L, Seno L, Valenza F, Valenzano A. A comprehensive approach to the automatic refinement and verification of access control policies. *Comput Secur* 2019;80:186–99. <http://dx.doi.org/10.1016/j.cose.2018.09.013>, URL: <https://www.sciencedirect.com/science/article/pii/S0167404818303870>.

- [58] Turkmen F, Crispo B. Performance evaluation of XACML PDP implementations. In: Proceedings of the 2008 ACM workshop on secure web services. 2008, p. 37–44.
- [59] Miseldine PL. Automated XACML policy reconfiguration for evaluation optimisation. In: Proceedings of the fourth international workshop on software engineering for secure systems. 2008, p. 1–8.
- [60] Marouf S, Shehab M, Squicciarini A, Sundareswaran S. Adaptive reordering and clustering-based framework for efficient XACML policy evaluation. *IEEE Trans Serv Comput* 2011;4(4):300–13. <http://dx.doi.org/10.1109/TSC.2010.28>.
- [61] Deng F, Yu Z, Liu W, Luo X, Fu Y, Qiang B, Xu C, Li Z. An efficient policy evaluation engine for XACML policy management. *Inform Sci* 2021;547:1105–21. <http://dx.doi.org/10.1016/j.ins.2020.08.044>, URL: <https://www.sciencedirect.com/science/article/pii/S0020025520308148>.
- [62] Narouei M, Takabi H. A nature-inspired framework for optimal mining of attribute-based access control policies. In: Lecture notes of the institute for computer sciences, social informatics and telecommunications engineering. Springer International Publishing; 2019, p. 489–506. http://dx.doi.org/10.1007/978-3-030-37231-6_29.
- [63] Hadj MAE, Benkaouz Y, Freisleben B, Erradi M. ABAC rule reduction via similarity computation. In: Networked systems. Springer International Publishing; 2017, p. 86–100. http://dx.doi.org/10.1007/978-3-319-59647-1_7.
- [64] Vaidya J, Atluri V, Guo Q, Adam N. Migrating to optimal RBAC with minimal perturbation. In: Proceedings of the 13th ACM symposium on access control models and technologies. SACMAT '08, New York, NY, USA: Association for Computing Machinery; 2008, p. 11–20. <http://dx.doi.org/10.1145/1377836.1377839>.
- [65] Jin C, Shen A, Yu W. The RBAC system based on role risk and user trust. *Int J Comput Commun Eng* 2016;5(5):374–80.
- [66] Dos Santos DR, Westphall CM, Westphall CB. A dynamic risk-based access control architecture for cloud computing. In: 2014 IEEE network operations and management symposium (NOMS). IEEE; 2014, p. 1–9.
- [67] Hadj MAE, Erradi M, Khoumsi A, Benkaouz Y. Validation and correction of large security policies: A clustering and access log based approach. In: 2018 IEEE international conference on big data (Big data). 2018, p. 5330–2. <http://dx.doi.org/10.1109/BigData.2018.8622610>.
- [68] Bauer L, Garriss S, Reiter MK. Detecting and resolving policy misconfigurations in access-control systems. *ACM Trans Inf Syst Secur* 2011;14(1). <http://dx.doi.org/10.1145/1952982.1952984>.
- [69] Hadj MAE, Khoumsi A, Benkaouz Y, Erradi M. Efficient security policy management using suspicious rules through access log analysis. In: Networked systems. Springer International Publishing; 2019, p. 250–66. http://dx.doi.org/10.1007/978-3-030-31277-0_16.
- [70] Hummer M, Kunz M, Netter M, Fuchs L, Pernul G. Adaptive identity and access management - contextual data based policies. *EURASIP J Inf Secur* 2016;2016(1). <http://dx.doi.org/10.1186/s13635-016-0043-2>.
- [71] Hill L. How automated access verification can help organizations demonstrate HIPAA compliance: A case study. *J Healthc Inf Manag* 2006;20(2):116–22.
- [72] Benedetti M, Mori M. Parametric RBAC maintenance via max-SAT. In: Proceedings of the 23rd ACM symposium on access control models and technologies. SACMAT '18, New York, NY, USA: Association for Computing Machinery; 2018, p. 15–25. <http://dx.doi.org/10.1145/3205977.3205987>.
- [73] Fuchs L, Pernul G. Reducing the risk of insider misuse by revising identity management and useraccount data. In: 2nd int. workshop on managing insider security threats, journal of wireless mobile networks, ubiquitous computing, and dependable applications (JoWUA). Morioka, Iwate, Japan; 2010, p. 14–28, URL: <https://epub.uni-regensburg.de/15129/>.
- [74] Colantonio A, Di Pietro R, Ocello A, Verde NV. Visual role mining: A picture is worth a thousand roles. *IEEE Trans Knowl Data Eng* 2012;24(6):1120–33. <http://dx.doi.org/10.1109/TKDE.2011.37>.
- [75] Guarnieri M, Arrigoni Neri M, Magri E, Mutti S. On the notion of redundancy in access control policies. In: Proceedings of the 18th ACM symposium on access control models and technologies. SACMAT '13, New York, NY, USA: Association for Computing Machinery; 2013, p. 161–72. <http://dx.doi.org/10.1145/2462410.2462426>.
- [76] Shamooni I, Rajpoot Q, Shibli A. Policy conflict management using XACML. In: 2012 8th international conference on computing and networking technology (INC, ICCIS and ICMI). 2012, p. 287–91.
- [77] Hu H, Ahn G-J, Kulkarni K. Discovery and resolution of anomalies in web access control policies. *IEEE Trans Dependable Secure Comput* 2013;10(6):341–54. <http://dx.doi.org/10.1109/tdsc.2013.18>.
- [78] Hounder F. Conflict detection and resolution of XACML policies (Master's thesis), University of Applied Sciences Rapperswil; 2010.
- [79] Deng F, Zhang L-Y. Elimination of policy conflict to improve the PDP evaluation performance. *J Netw Comput Appl* 2017;80:45–57. <http://dx.doi.org/10.1016/j.jnca.2016.12.001>.
- [80] Dia OA, Farkas C. A practical framework for policy composition and conflict resolution. *Int J Secure Softw Eng* 2012;3(4):1–26. <http://dx.doi.org/10.4018/jsse.2012100101>.
- [81] Currey J, McKinstry R, Dadgar A, Gritter M. Informed privilege-complexity trade-offs in RBAC configuration. In: Proceedings of the 25th ACM symposium on access control models and technologies. SACMAT '20, New York, NY, USA: Association for Computing Machinery; 2020, p. 119–30. <http://dx.doi.org/10.1145/3381991.3395597>.
- [82] Hu J, Zhang Y, Li R. Towards automatic update of access control policy. In: Proceedings of the 24th international conference on large installation system administration. LISA'10, USA: USENIX Association; 2010, p. 1–7.
- [83] Argento L, Margheri A, Paci F, Sassone V, Zannone N. Towards adaptive access control. In: Data and applications security and privacy XXXII. Springer International Publishing; 2018, p. 99–109. http://dx.doi.org/10.1007/978-3-319-95729-6_7.
- [84] Benkaouz Y, Erradi M, Freisleben B. Work in progress: K-nearest neighbors techniques for ABAC policies clustering. In: Proceedings of the 2016 ACM international workshop on attribute based access control. ABAC '16, New York, NY, USA: Association for Computing Machinery; 2016, p. 72–5. <http://dx.doi.org/10.1145/2875491.2875497>.
- [85] Hadj MAE, Khoumsi A, Benkaouz Y, Erradi M. Formal approach to detect and resolve anomalies while clustering ABAC policies. *ICST Trans Secur Saf* 2018;5(16):156003. <http://dx.doi.org/10.4108/eai.13-7-2018.156003>.
- [86] Hein P, Biswas D, Martucci LA, Muhlhäuser M. Conflict detection and lifecycle management for access control in publish/subscribe systems. In: 2011 IEEE 13th international symposium on high-assurance systems engineering. 2011, p. 104–11. <http://dx.doi.org/10.1109/HASE.2011.50>.
- [87] Hu H, Ahn G-J, Kulkarni K. Detecting and resolving firewall policy anomalies. *IEEE Trans Dependable Secure Comput* 2012;9(3):318–31. <http://dx.doi.org/10.1109/tdsc.2012.20>.
- [88] Oberholzer S. Optimizing XACML policies. University of Applied Sciences Rapperswil; 2011.
- [89] Stepien B, Matwin S, Felty A. An algorithm for compression of XACML access control policy sets by recursive subsumption. In: 2012 seventh international conference on availability, reliability and security. 2012, p. 161–7. <http://dx.doi.org/10.1109/ARES.2012.38>.
- [90] Al-Kahtani M, Sandhu R. A model for attribute-based user-role assignment. In: 18th annual computer security applications conference, 2002. Proceedings. 2002, p. 353–62. <http://dx.doi.org/10.1109/CSAC.2002.1176307>.
- [91] Al-Kahtani M, Sandhu R. Induced role hierarchies with attribute-based RBAC. In: Proceedings of the eighth ACM symposium on access control models and technologies. SACMAT '03, New York, NY, USA: Association for Computing Machinery; 2003, p. 142–8. <http://dx.doi.org/10.1145/775412.775430>.
- [92] Baumgrass A. Deriving current state RBAC models from event logs. In: 2011 sixth international conference on availability, reliability and security. 2011, p. 667–72. <http://dx.doi.org/10.1109/ARES.2011.104>.
- [93] Chakraborty S, Ray I. TrustBAC: Integrating trust relationships into the RBAC model for access control in open systems. In: Proceedings of the eleventh ACM symposium on access control models and technologies. SACMAT '06, New York, NY, USA: Association for Computing Machinery; 2006, p. 49–58. <http://dx.doi.org/10.1145/1133058.1133067>.
- [94] Gal-Oz N, Gonen Y, Yahalom R, Gudes E, Rozenberg B, Shmueli E. Mining roles from web application usage patterns. In: Trust, privacy and security in digital business. Springer Berlin Heidelberg; 2011, p. 125–37. http://dx.doi.org/10.1007/978-3-642-22890-2_11.
- [95] Guo Q, Vaidya J, Atluri V. The role hierarchy mining problem: Discovery of optimal role hierarchies. In: 2008 annual computer security applications conference (ACSAC). 2008, p. 237–46. <http://dx.doi.org/10.1109/ACSAC.2008.38>.
- [96] Han D-j, Zhuo H-k, Xia L-t, Li L. Permission and role automatic assigning of user in role-based access control. *J Central South Univ* 2012;19(4):1049–56. <http://dx.doi.org/10.1007/s11771-012-1108-0>.
- [97] Herzberg A, Mass Y, Mihaeli J, Naor D, Ravid Y. Access control meets public key infrastructure, or: assigning roles to strangers. In: Proceeding 2000 IEEE symposium on security and privacy. S P 2000. 2000, p. 2–14. <http://dx.doi.org/10.1109/SECPRI.2000.848442>.
- [98] Hu J, Zhang Y, Li R, Lu Z. Role updating for assignments. In: Proceedings of the 15th ACM symposium on access control models and technologies. SACMAT '10, New York, NY, USA: Association for Computing Machinery; 2010, p. 89–98. <http://dx.doi.org/10.1145/1809842.1809859>.
- [99] Hu J, Khan KM, Zhang Y, Bai Y, Li R. Role updating in information systems using model checking. *Knowl Inf Syst* 2016;51(1):187–234. <http://dx.doi.org/10.1007/s10115-016-0974-4>.
- [100] Huang J, Nicol DM, Bobba R, Huh JH. A framework integrating attribute-based policies into role-based access control. In: Proceedings of the 17th ACM symposium on access control models and technologies. SACMAT '12, New York, NY, USA: Association for Computing Machinery; 2012, p. 187–96. <http://dx.doi.org/10.1145/2295136.2295170>.

- [101] Lu J, Xu D, Jin L, Han J, Peng H. On the complexity of role updating feasibility problem in RBAC. *Inform Process Lett* 2014;114(11):597–602. <http://dx.doi.org/10.1016/j.ipl.2014.06.003>.
- [102] Lu J, Xin Y, Peng H, Han J, Lin F. Supporting user authorization queries in RBAC systems by role-permission reassignment. In: *Cyberspace safety and security*. Springer International Publishing; 2017, p. 468–76. http://dx.doi.org/10.1007/978-3-319-69471-9_35.
- [103] Ni Q, Lobo J, Calo S, Rohatgi P, Bertino E. Automating role-based provisioning by learning from examples. In: *Proceedings of the 14th ACM symposium on access control models and technologies*. SACMAT '09, New York, NY, USA: Association for Computing Machinery; 2009, p. 75–84. <http://dx.doi.org/10.1145/1542207.1542222>.
- [104] Pang C, Hansen D, Maeder A. Managing RBAC states with transitive relations. In: *Proceedings of the 2nd ACM symposium on information, computer and communications security*. ASIACCS '07, New York, NY, USA: Association for Computing Machinery; 2007, p. 139–48. <http://dx.doi.org/10.1145/1229285.1229306>.
- [105] Pang C, Zhang X, Zhang Y, Ramamohanarao K. The efficient maintenance of access roles with role hiding. In: Das G, Sarda NL, Reddy PK, editors. *Proceedings of the 14th international conference on management of data*, December 17–19, 2008, IIT Bombay, Mumbai, India. Computer Society of India / Allied Publishers; 2008, p. 139–49, URL: <http://www.cse.iitb.ac.in/%7Ecomad/2008/PDFs/13.pdf>.
- [106] Rao KR, Nayak A, Ray IG, Rahulamathavan Y, Rajarajan M. Role recommender-RBAC: Optimizing user-role assignments in RBAC. *Comput Commun* 2021;166:140–53. <http://dx.doi.org/10.1016/j.comcom.2020.12.006>.
- [107] Saffarian M, Tang Q, Jonker W, Hartel P. *Dynamic user role assignment in remote access control*. CTIT technical report series, TR-CTIT-09-14, Netherlands: Centre for Telematics and Information Technology (CTIT); 2009, eemcs-eprint-15311.
- [108] Shañiq B, Vaidya JS, Ghafoor A, Bertino E. A framework for verification and optimal reconfiguration of event-driven role based access control policies. In: *Proceedings of the 17th ACM symposium on access control models and technologies*. SACMAT '12, New York, NY, USA: Association for Computing Machinery; 2012, p. 197–208. <http://dx.doi.org/10.1145/2295136.2295172>.
- [109] Sheng S, Osborn SL. A classifier-based approach to user-role assignment for web applications. In: *Lecture notes in computer science*. Springer Berlin Heidelberg; 2004, p. 163–71. http://dx.doi.org/10.1007/978-3-540-30073-1_12.
- [110] Strembeck M. A role engineering tool for role-based access control. In: *Proceedings of the third symposium on requirements engineering for information security* SREIS. 2005, p. 1–8.
- [111] Strembeck M. Scenario-driven role engineering. *IEEE Secur Priv Mag* 2010;8(1):28–35. <http://dx.doi.org/10.1109/msp.2010.46>.
- [112] Takabi H, Amini M, Jalili R. Trust-based user-role assignment in role-based access control. In: *2007 IEEE/ACS international conference on computer systems and applications*. 2007, p. 807–14. <http://dx.doi.org/10.1109/AICCSA.2007.370725>.
- [113] Takabi H, Joshi JB. StateMiner: An efficient similarity-based approach for optimal mining of role hierarchy. In: *Proceedings of the 15th ACM symposium on access control models and technologies*. SACMAT '10, New York, NY, USA: Association for Computing Machinery; 2010, p. 55–64. <http://dx.doi.org/10.1145/1809842.1809853>.
- [114] Yi-qun Z, Jian-hua L, Quan-hai Z. A general attribute based RBAC model for web service. In: *IEEE international conference on services computing (SCC 2007)*. 2007, p. 236–9. <http://dx.doi.org/10.1109/SCC.2007.8>.
- [115] Zhang D, Ramamohanarao K, Ebringer T. Role engineering using graph optimisation. In: *Proceedings of the 12th ACM symposium on access control models and technologies*. SACMAT '07, New York, NY, USA: Association for Computing Machinery; 2007, p. 139–44. <http://dx.doi.org/10.1145/1266840.1266862>.
- [116] Zhang W, Chen Y, Gunter C, Liebovitz D, Malin B. Evolving role definitions through permission invocation patterns. In: *Proceedings of the 18th ACM symposium on access control models and technologies*. SACMAT '13, New York, NY, USA: Association for Computing Machinery; 2013, p. 37–48. <http://dx.doi.org/10.1145/2462410.2462422>.
- [117] Zong Y, Bhargava B, Mahoui M, Zhong Y. Trustworthiness based authorization on WWW. In: *IEEE workshop on security in distributed data warehousing*. 2011, p. 1–6.
- [118] Hummer M, Groll S, Kunz M, Fuchs L, Pernul G. Measuring identity and access management performance - an expert survey on possible performance indicators. In: *Proceedings of the 4th international conference on information systems security and privacy*. SCITEPRESS - Science and Technology Publications; 2018, p. 233–40. <http://dx.doi.org/10.5220/0006557702330240>.
- [119] Schrimpff A, Drechsler A, Dagianis K. Assessing identity and access management process maturity: First insights from the german financial sector. *Inf Syst Manage* 2021;38(2):94–115. <http://dx.doi.org/10.1080/10580530.2020.1738601>, [arXiv:https://doi.org/10.1080/10580530.2020.1738601](https://doi.org/10.1080/10580530.2020.1738601).
- [120] Ahn G-J, Sandhu R. Role-based authorization constraints specification. *ACM Trans Inf Syst Secur* 2000;3(4):207–26. <http://dx.doi.org/10.1145/382912.382913>.
- [121] Fuchs L, Pernul G. Supporting compliant and secure user handling - a structured approach for in-house identity management. In: *The second international conference on availability, reliability and security (ARES'07)*. 2007, p. 374–84. <http://dx.doi.org/10.1109/ARES.2007.145>.
- [122] Fuchs L, Broser C, Pernul G. Different approaches to in-house identity management - justification of an assumption. In: *2009 international conference on availability, reliability and security*. 2009, p. 122–9. <http://dx.doi.org/10.1109/ARES.2009.154>.
- [123] Hornsteiner M, Groll S, Puchta A. Towards a user-centric IAM entitlement shop - learnings from the e-commerce. In: *13th international conference on security of information and networks*. SIN 2020, New York, NY, USA: Association for Computing Machinery; 2020, p. 1–4. <http://dx.doi.org/10.1145/3433174.3433585>.
- [124] Royer D. *Enterprise identity management*. In: Fischer-Hübner S, Duquenoey P, Zuccato A, Martucci L, editors. *The future of identity in the information society*. Boston, MA: Springer US; 2008, p. 433–46.
- [125] Bobba R, Gavrilas S, Gligor V, Khurana H, Koleva R. Administering access control in dynamic coalitions. In: *Proceedings of the 19th conference on large installation system administration conference - Volume 19*. LISA '05, USA: USENIX Association; 2005, p. 23.
- [126] Heinrich B, Klier M. A novel data quality metric for timeliness considering supplemental data. In: *Newell S, Whitley EA, Pouloudi N, Wareham J, Mathiasen L, editors. 17th European conference on information systems, ECIS 2009, Verona, Italy, 2009*. 2009, p. 2651–62, URL: <http://aisel.aisnet.org/ecis2009/14>.
- [127] Menges F, Latzo T, Vielberth M, Sobola S, Pöhls HC, Taubmann B, Köstler J, Puchta A, Freiling F, Reiser HP, et al. Towards GDPR-compliant data processing in modern SIEM systems. *Comput Secur* 2021;103:102165.
- [128] Das S, Sural S, Vaidya J, Atluri V. HyPE: A hybrid approach toward policy engineering in attribute-based access control. *IEEE Lett Comput Soc* 2018;1(2):25–9. <http://dx.doi.org/10.1109/locs.2018.2889980>.



Sascha Kern studied Management Information Systems at the University of Cologne and the University of Regensburg with a focus on IT security. Since 2016, he works as a software engineer at Nexis GmbH and develops solutions for the implementation of IAM measures in coordination with customers and IAM experts. During research work, Sascha Kern examines how the quality of data can be assessed and improved in the context of IAM to enable a sufficient level of IT security.



Thomas Baumer studied Management Information Systems with a specialization in IT security at the University of Regensburg and KU Leuven. He graduated within the Honors Elite Program as M.Sc. with Honors at Regensburg in 2020. Since then he is working as a software engineer at Nexis GmbH, a spin-off from the Chair of Information Systems I (Prof. Dr. Pernul, University of Regensburg) specializing in Identity and Access Governance & Analytics. There he pursues his Ph.D. with a research focus in maintaining a high and usable IT security level in a changing environment by grasping synergies from research and practice.



Sebastian Groll studied Business Informatics within the Honors Elite Program at University of Regensburg and at Kingston University in London. His major field of study was IT Security in his master degree course. Since January 2018, Sebastian Groll is a research assistant at the Chair of Information Systems I (Prof. Dr. Pernul). His research is supplemented by his practical experience within Identity and Access Management as he is parallelly working as a technical consultant and developer at the Nexis GmbH. Nexis is specialized on Identity and Access Governance & Analytics and offers the software solution Nexis Controle and consulting services.

S. Kern et al.



Ludwig Fuchs studied Information Systems at the University of Regensburg, Germany and completed his dissertation in 2009. He studied and researched at the University of York (UK) and the University of Texas (San Antonio, USA) together with well-known academics in the field of IT security (e.g. Prof. Dr. Ravi Sandhu, “RBAC”). His research interest comprises Identity Management within mid-sized and large organizations. Over the last ten years, Ludwig Fuchs gathered practical and academic experience and published at several IT security conferences. His expert knowledge has been underlined throughout his work in industry projects, bridging the gap between research and practice.

Journal of Information Security and Applications 70 (2022) 103301



Günther Pernul (Member, IEEE) received the diploma and Ph.D. degrees (Hons.) in business informatics from the University of Vienna, Austria. He is currently a Professor with the Department of Information Systems, University of Regensburg, Germany. Previously, he held positions at the University of Duisburg–Essen, Germany; the University of Vienna; the University of Florida, Gainesville; and the College of Computing, Georgia Institute of Technology, Atlanta. His research interests include data and information-security aspects, data protection and privacy, data analytics, and advanced datacentric applications.

2 Identity and Access Management Metrics

Current status:	Under Review
Journal:	Submitted to: ACM Computing Surveys (CSUR)
Date of acceptance:	-
Full citation:	Thomas Baumer, Sascha Kern, Ludwig Fuchs, and Günther Pernul. Identity and Access Management Metrics. Submitted to: <i>ACM Computing Surveys</i> .
Authors contributions:	Thomas Baumer 50% Sascha Kern 35% Ludwig Fuchs 5% Günther Pernul 10%

Journal Description: These comprehensive, readable surveys and tutorial papers give guided tours through the literature and explain topics to those who seek to learn the basics of areas outside their specialties in an accessible way. The carefully planned and presented introductions in Computing Surveys (CSUR) are also an excellent way for researchers and professionals to develop perspectives on, and identify trends in complex technologies. Contributions which bridge existing and emerging technologies (such as machine learning) with a variety of science and engineering domains in a novel and interesting way are also welcomed.

Identity and Access Management Metrics

THOMAS BAUMER, SASCHA KERN, and LUDWIG FUCHS, Nexis GmbH, Germany

GÜNTHER PERNUL, Universität Regensburg, Germany

Identity and Access Management (IAM) is an interdisciplinary challenge for organizations that requires carefully designed business processes, supporting technologies, and effective access control policies. Its goals encompass security, compliance, operations, and high quality. IAM metrics aid these goals' strategic steering and alignment. This work contributes with an overview of crucial IAM metrics by reviewing the present body of literature, connecting them with IAM goals and audiences, and grouping them for seven IAM-relevant perspectives to foster comprehension.

CCS Concepts: • **General and reference** → **Surveys and overviews; Metrics**; • **Security and privacy** → **Security services**; • **Information systems** → **Enterprise information systems**; • **Applied computing** → **Business-IT alignment; IT governance**.

Additional Key Words and Phrases: Identity and Access Management, Metrics, Measurements, Governance, Data Quality

ACM Reference Format:

Thomas Baumer, Sascha Kern, Ludwig Fuchs, and Günther Pernul. 2023. Identity and Access Management Metrics. 1, 1 (December 2023), 35 pages. <https://doi.org/XXXXXXXX.XXXXXXX>

1 INTRODUCTION

Organizations encounter vast difficulties securing and efficiently organizing users' access to digital resources. While mature technologies for enforcing authentication and authorization are available, managing these security mechanisms remains challenging. Numerous threat reports thus list insufficient or broken access control among the most frequent and critical Information Technology (IT) security vulnerabilities [2, 16, 94]. Such vulnerabilities can open the door for internal and external attackers, bypassing authentication mechanisms or using excessive privileges for malicious action. Organizations worldwide counter this threat by investing significant shares of their IT security budgets into Identity and Access Management (IAM) measures [47]. IAM metrics aid the management of these measures as they help to identify the need for action and to assess their effectiveness. While IAM and related problems such as *identity chaos* [45] and the *policy explosion problem* [36] have received broad research attention, to the best of the authors' knowledge, no scientific study exists that holistically examines IAM metrics and their impact on IAM goals. This work presents a literature survey identifying 43 strategic IAM metrics based on 7 perspectives derived from strategic IAM goals. We summarize the metrics for each perspective and categorize them by their target, goal interdependence, and audience. We then discuss insights based on the perspectives, the relevance of IAM goals, and understandability.

The remaining paper outlines as follows. Section 2 presents the background for IAM and metrics. With the method of Section 3, Section 4 derives assumed IAM goals and audiences for this work. Sections 5 and 6 then enumerate, categorize, and discuss IAM metrics. Section 7 aligns our contribution to related work. Lastly, Section 8 concludes.

Authors' addresses: Thomas Baumer, Thomas.Baumer@nexis-secure.com; Sascha Kern; Ludwig Fuchs, Nexis GmbH, Rudolf-Vogt-Straße 6, 93053, Regensburg, Germany; Günther Pernul, Universität Regensburg, Universitätsstraße 31, 93053, Regensburg, Germany.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2023 Association for Computing Machinery.

Manuscript submitted to ACM

Manuscript submitted to ACM

1

2 BACKGROUND AND TERMINOLOGY DISCUSSION

2.1 Identity and Access Management

IAM is a domain of IT management dealing with the administration of (digital) identities and the secure access of users to digital resources within an organizational context. It manages a range of processes, policies, and technologies designed to ensure that IT security objectives of authentication and authorization are met without disproportionately limiting operations. The tasks required for this include the creation and management of user accounts, the provision, and verification of user credentials, the definition and enforcement of adequate authorizations, the secure and privacy-preserving management of related data, the assurance of regulatory compliance, as well as a large number of subordinate or ancillary activities. In addition, IAM has to interact with several business processes that do not originate from IAM: Prominent examples are a person joining the organization, changing department affiliation, or leaving the organization (also known as joiner, mover and leaver processes); or changes in the IT infrastructure that require an adjustment of the access structures, for example when a new application system is launched. The execution of these activities is often supported by specific technologies, such as directories or dedicated IAM and Identity Governance and Administration (IGA) systems, which are specialized in the processing of IAM-specific data and processes. [42, 112, 138]

Since IAM controls (like Principle of Least Privilege (PoLP)) are a basis for providing secure user access, they are demanded by regulative frameworks, such as the United States of America (USA) Sarbanes-Oxley Act (SOX) [149], the European Union (EU) General Data Protection Regulation (GDPR) [37], USA Health Insurance Portability and Accountability Act (HIPAA) [148], the Basel III accords [7] or the International Organization for Standardization (ISO) 27001 standard [77]. To comply with these regulations, organizations must prove their compliance during external audits. The fields of action involved in IAM result in several actors with different goals and responsibilities, which may vary from one organization to another. Section 4 details these and their connection to IAM metrics.

2.2 Metrics and Measurements

The value of scientific work is determined by applying evaluation measurements. However, research differs in scope, methodologies, and publication quality, leading to confusion about terms even within a single research discipline. Besides inconsistencies in the definitions of metrics and measurements, other terms like Key Performance Indicators (KPIs), scores, assessments, indicators, etc. are also used. Some efforts seek to standardize metrics and measurements for IT security and data quality. These terms are introduced in this section and apply to this work. [18, 23, 61, 165]

Most confusion about metrics and measurement arises because of their similarity. Both concepts take numeric values as input and output, leading to similar behavior in collecting data, addressing the audience, visualization, or reporting. However, both concepts differ in their abstraction levels [18]. While a measurement objectively measures evidence, a metric (somewhat) subjectively interprets a set of measurements on a higher abstraction level. An example of a measurement is the password length which is objectively measured by counting password characters. While this measurement is a standalone property, its actual expressiveness is limited since it could not distinguish a password like "11111111" from a stronger one. In an intuitive process, one might combine the password length with additional measurements like the used charsets or its monotonicity to raise expressiveness. Measurements are converted into a more expressive (password strength) metric using a formula that determines the selection and weighting of considered measurements. While these formulas are subjective and context-specific, the metric concept *password strength* remains objective. While this aggregation raises expressiveness, the standalone measurements preserve as drill-down functionalities, allowing decomposing metrics for deeper insights. Furthermore, some fine-tuning of a metric to fit the subjective

Manuscript submitted to ACM

context is required to maximize its expressiveness. E.g., a dedicated, standalone study [31] compares different formulas for password strength metrics of popular web services. It shows the depth of analyzing even one metric in detail. In contrast, our study focuses on an overview of IAM metrics. It does not detail specific measurements and formulas that might support a metric. Figure 1 visualizes measurements, metrics, and their relationships. We let measurements and metrics inherit from an abstract *measurable* entity for shared properties. The following paragraphs detail these entities.




Fig. 1. Metrics and related entities (own depiction, derived from Black et al. [18] and Chew et al. [23]).

Measurables hold shared properties [23] of measurements and metrics. A measurable requires a unique *identifier*, enabling structured processing. Another property is *goals*. A metric might affect goals positively and negatively. For example, a good evaluation of a password strength metric might support a security goal while hampering operations [67]. The *target* property defines the measurement's desired value (like min or max). The *audience* property specifies interest groups for measurables. A measurable has a *description* property for readability purposes. Another property is *frequency*, distinguishing between collection and reporting frequency. The frequency property indicates *history*, storing historic values scheduled by the collection frequency. The next property is *automation*. Automation is a best practice for measurables retrievable from databases. However, automation might not be feasible for data collection based on interviews, user studies, etc. Finally, the *reporting format* property indicates the desired visualization.

Measurements inherit properties described for the measurable entity. Additional properties are evidence, unit, and data source. The *evidence* exactly describes the data collection. It can refer to a database query, a manual user study, etc. The *unit* property indicates the measurement unit. Possible values are percentage, number, etc. Finally, the *data source* property defines the data location for retrieval. It includes databases, tools, organizations, specific roles, or users.

Formula is a relationship between metrics and measurements. It thus contains a set of relationships to *measurements* with their *weighting* and *combination rule* for the metric (e.g., by [15]).

Metrics inherit measurables' properties and are this study's primary focus. Metrics are calculated using their relationship property *formula*. Additional requirements are defined for IT security [165] and data quality [61]. A metric thus needs to be *bounded*. One metric bound represents the perfectly good (IT security, data quality, etc.), while the other represents the perfectly poor [61]. The *scale* of a metric is metrical, which includes ratio and interval scaling. That allows for the comparison of metric values. Furthermore, metrics satisfy *quality* based on objectivity, reliability, and validity. (i) Objectivity describes the independence of external influences. E.g., the objectivity of a password strength metric is limited if relevant measurements like the password length or the monotonicity are not covered. (ii) Reliability ensures the reproducibility of determining a metric value. E.g., the value of a password strength metric should not change when applied *ceteris paribus*. (iii) Validity means that a metric actually measures the theoretical construct it tries to. E.g., if a password strength metric also considers an unrelated user satisfaction, its validity decreases. Another requirement is *sound aggregation*. A metric is thus consistently applicable for single and multiple data values (e.g., for a single user or a whole department). Finally, a metric has to be *economically efficient*, meaning the costs of defining and (automatically [165]) measuring the metrics cannot outweigh the expected payoff.

4

Baumer et al.

3 METHOD

The method for this study follows the guidelines of Levy and Ellis [87]. Figure 2 visualizes the general approach for this work. Our primary goal is to provide a broad yet understandable and quantitative overview of Identity and Access Management (IAM) within organizations using metrics. Therefore, we studied general terminologies and concepts for metrics and IAM in Section 2. These essentials highlight the relevance of goals and audiences for an effective alignment of metrics and strategy. Section 4 examines both in greater detail for IAM.




Fig. 2. General approach for this study.

Our literature analysis targets publications with quantitative assessments of IAM concepts. The search utilizes high-quality online libraries, like Springer Link¹, AIS eLibrary², IEEE Xplore³, ScienceDirect⁴, ACM Digital Library⁵, and Google Scholar⁶, with combinations of search terms like IAM, IGA, Identity Management (IdM) and metrics, measurements, assessments, etc. Furthermore, a forward and backward search refined the results. We identified 52 contributions with sufficient scientific quality and suitable scope. The results got refined by subsequent synthesis steps.

Next, we extracted and synthesized IAM metrics with a divide and conquer approach, operating top-down and bottom-up. First, we assumed various IAM perspectives. These perspectives align with IAM processes, like the identity and policy lifecycle or governance [112], and IAM quality goals for identities, policies, processes, and software [42]. This allows for an improved investigation of each perspective, either top-down or bottom-up. For the top-down ones, we applied present metric sets and connect them to IAM. We found and applied these metric sets for identities, policies, processes, and software quality. For the bottom-up ones, we aggregated IAM measurements and metrics found in the literature for the remaining perspectives: identity lifecycle, policy lifecycle, governance. A mind map helped us to outline the aggregation by matching the found IAM measurements and metrics. Additionally, for understandability purposes, we also restricted ourselves to a limited amount of metrics for each perspective. Renowned psychological research shows that presenting too many concepts simultaneously leads to comprehension issues due to the limits of human information processing. Presenting at max. 7 ± 2 [96] or even better only 4 items [29] is more understandable.

In summary, we synthesized 43 IAM metrics spanning over 7 IAM-relevant perspectives in Section 5. Finally, Section 6 discusses implications of the IAM metrics and their perspectives, alignment to IAM goals, and understandability.

4 GOALS AND AUDIENCES FOR IDENTITY AND ACCESS MANAGEMENT METRICS

The performance of an organization is generally determined by the alignment of cost reduction, revenue growth, and quality improvement [110]. We use these three dimensions to structure strategic IAM goals. Cutting unnecessary operational costs, like expenses, efforts, or resources, achieves cost reduction. Delivering value-adding products and services and satisfying actual needs concludes in revenue growth. Quality improvement helps cost reduction and revenue growth by fostering refined, less error-prone processes. IAM aligns with these strategic dimensions by improving

¹<https://link.springer.com/advanced-search>

²<https://aisel.aisnet.org/do/search/advanced/>

³<https://ieeexplore.ieee.org/search/advanced>

⁴<https://www.sciencedirect.com/search/entry>

⁵<https://dl.acm.org/search/advanced>

⁶<https://scholar.google.de/>

Manuscript submitted to ACM

quality and reducing costs. However, value addition usually remains unaffected since IAM supports an organization internally while not improving the delivered products or services (except IAM itself is the generated value by the organization, cmp. Customer Identity and Access Management (CIAM)).

Cost reduction by IAM distinguishes security, compliance, and operations. (i) Security focuses on preventing unauthorized access to valuable resources like trade secrets or customer data. These leaks might jeopardize an organization's trade secrets or damage its reputation. IAM provides basic mechanisms like password management, authentication, authorization, and access logging. (ii) Compliance requirements are imposed by (inter-) national authorities and regulate access management and audit essentials. Penalties apply for violations. For large organizations, avoiding compliance violations and their penalties frequently represents a main driver for IAM capabilities. (iii) Operations aim for smooth processes. Easy and fast access to required resources enables efficient task processing. Missing permissions disrupt important user tasks, which wastes the time of single users or puts whole organizations on hold. For example, as a tool, Single Sign On (SSO) software improves operations by easing organization-wide user authentication.

Quality-based goals support cost-based and value-based goals. For this work, we break down IAM quality into further subgoals: IAM processes, identities, IAM policies, and IAM software. These align closely with the domains described by [42] and add identity quality as a further goal. (i) IAM utilizes processes that require high quality in their definition and execution to ensure the desired outcome. These IAM processes span from governance to modeling processes, like identity and policy lifecycles, access reviews, or operational processes (e.g., authentication, self-service, or help desk processes). (ii) The quality of both identities and IAM policies closely tie to the quality of IAM data: Since digital identities are commonly defined as a set of attributes, their quality reduces to the quality of the comprised attributes [116]. (iii) The quality of IAM policies is more complex, as the underlying access control models require consideration [84]. While many authors conceptualize IAM policy quality as data quality with dimensions such as accuracy, complexity, and redundancy, other authors (especially in the eXtensible Access Control Markup Language (XACML) domain) conceptualize it as software quality and focus on features such as evaluation runtime or testing of policy configurations [83, 85]. (iv) Finally, IAM software quality is the foundation for IAM effectiveness since the other quality considerations depend on high-quality software. Additionally, the quality of these components stands in mutual connection with the quality of other elements that are not directly related to IAM, such as data quality, process quality, or information system quality.




Fig. 3. Strategic IAM goals (own depiction, derived from Oh and Pinsonneault [110], Hummer et al. [67] and Fuchs et al. [42]).

Figure 3 lists cost-based (left) and quality-based (right) IAM goals. These goals are drivers that shape IAM strategy [42]. IAM metrics assess the degree of achievement of these IAM goals. They are relevant tools for identifying problems (e.g., outdated identity attributes), selecting and prioritizing investments, and determining the success of IAM efforts.

The audience determines for whom an IAM metric is relevant. In real-world organizations, many actors are involved in or affected by IAM processes. However, neither the job description (e.g., recruiter) nor the organizational affiliation (e.g., department head) suits well to describe a role in an IAM context since these are organization-specific. Instead, we focus on prototypical stakeholder roles to determine the audience of an IAM metric. Depending on their tasks

6

Baumer et al.

261 and responsibilities, a person holds one, several, or none of these stakeholder roles. Despite some efforts [23, 42, 66],
 262 IAM stakeholders differ in scope and detail across research. This work derives five audience roles: (i) A governance
 263 audience contains management, compliance, and audit responsibilities. (ii) Identity admins, like supervisors or Human
 264 Resources (HR) employees, are responsible for keeping identities up-to-date across all lifecycle phases. (iii) Policy
 265 admins design and maintain policies involving policy engineers or owners. (iv) Operations keep IAM systems running,
 266 spanning over help desks, DevOps, or provisioning engineers. (v) Lastly, users need access to fulfill their daily duties.
 267
 268

269 **5 IDENTITY AND ACCESS MANAGEMENT METRICS**
 270

271 This Section 5 presents the resulting IAM metric set based on the literature survey. We organize these metrics by
 272 IAM-relevant perspectives [112, 138]. On the one hand, perspectives derived from the cost-reduction goal comprise the
 273 primary IAM process frameworks: the identity lifecycle (Section 5.1), policy lifecycle (Section 5.2), and a management
 274 and governance perspective (Section 5.3). On the other hand, we map a perspective for each quality goal, namely
 275 identity quality (Section 5.4), process quality (Section 5.5), policy quality (Section 5.5, and software quality (Section 5.7).
 276

277 Each perspective also summarizes the IAM metrics within the Tables 1, 2, 3, 4, 5, 6, and 7. The Tables include
 278 the metric name and a target determining whether the metric should strive to be minimized, maximized, or target a
 279 context-specific value. We also depict the metrics' relationship towards IAM goals and audience. For the IAM goals, we
 280 track whether the metric has a positive (+) or negative (-) impact. The goals are abbreviated⁷ for a compact presentation.
 281 For the IAM audience, we also abbreviate⁸ the prototypical audiences and tick the relevant ones for each metric.
 282
 283

284 **5.1 Identity Lifecycle Perspective**
 285

286 The identity lifecycle defines creating, managing, and deleting digital identities and links to joiner, mover, leaver,
 287 or support processes [75, 112]. The lifecycle accompanies an identity in all phases and ensures granting required
 288 but not excessive access. The lifecycle includes a timely (de-)provisioning of identities [10], education for members
 289 of an organization, and revoking excessive access. Identities also need operational support, like when requesting
 290 missing Access Control Policies (ACPs) or changing credentials. Since authentication schemes use knowledge-based,
 291 biometric-based, ownership-based, and combinations as multi-factor-based authentication factors, these strengths are
 292 measurable [109]. This lifecycle is relevant for security and operations goals. Missing authorizations impede processes,
 293 while excessive authorizations violate the PoLP. Table 1 summarizes the metrics for this perspective. [42, 125]
 294
 295

296
 297 Table 1. Identity and Access Management Metrics: Identity Lifecycle Perspective (own depiction based on a bottom-up derivation).
 298

299

Metric	Target	Cost Goal				Quality Goal				IAM Audience				
		COST	SEC	COM	OPS	PRC	ID	POL	SW	GOV	ID	POL	OPS	USR
Readiness	max	-	-	+	+		+		+	✓	✓	✓	✓	✓
Education	target	-	+	-	+				+	✓	✓	✓	✓	✓
Orphans	min	+	-	-	+	+	-			✓	✓	✓	✓	✓
Excess	target	-	-	-	-	-	-			✓	✓	✓	✓	✓
Password Strength	max	-	+	+	-		+		+	✓	✓	✓	✓	✓
Biometric Strength	max	-	+	+	-		+		+	✓	✓	✓	✓	✓
Ownership Strength	max	-	+	+	-		+		+	✓	✓	✓	✓	✓
Multi-factor Strength	max	-	+	+	-		+		+	✓	✓	✓	✓	✓

300
301
302
303
304
305
306
307

308
 309 ⁷Goal abbreviations: cost reduction (COST), security (SEC), compliance (COM), operations (OPS), process quality (PRC), identity quality (ID), policy
 310 quality (POL), and software quality (SW).
 311 ⁸Audience abbreviations: management and governance (GOV), identity admins (ID), policy admins (POL), operations (OPS), and the end users (USR).
 312

Manuscript submitted to ACM

Readiness describes the degree to which identities are ready and prepared to operate. It includes creating identities and accounts with valid appropriate authorizations and avoiding uncalled down-times for users. Joiner and mover processes of the identity lifecycle or provisioning tasks are thus especially relevant. Suitable measurements are the duration upon readiness achievement, error rates for identity, and account creation or costs [15, 67, 135]. Despite the technical dimension, readiness might also encompass the IAM education of the organization's members, considered in a separate metric. Readiness eases operations and supports compliance by indicating controls for timely authorization grants. However, achieving readiness either requires cost-intensive manual administration or preferably providing an appropriate automatism. As quality side-effects, readiness mainly affects identities and software. Timely propagated identities and accounts benefit the identity quality across multiple applications. Accessible and high-quality IAM software also encompasses means for external administration of their identities. Relevant IAM audiences cover governance, identity admins, operations, and users. Identity admins and users are directly affected by this metric. For operations, lower readiness evaluations signal more prospective support. E.g., expecting more requests for missing access. Finally, a governance audience concerns readiness as an indicator for good IAM controls.

Education expresses IAM awareness and knowledge. Even the best IAM systems render useless for arbitrary authorization grants or authorized yet inappropriate user behavior. All users that require sensible authorizations need at least a basic IAM and IT security training. General user guidance and awareness on *Dos and Don'ts* should be accessible and taught. This guidance includes appropriate password handling [151], recognizing and withdrawing from suspicious situations (e.g., *say no*), or awareness of sharing protected information on different media (e.g., internal communication, on-premise, email, etc.) [160]. More vulnerable groups also require more intensive training. Vulnerable groups include front liners like help-desk or telephone operators and privileged users. These groups should not hesitate to escalate or report incidents on recognition boldly [160]. The *dual knowledge gap* [40] points out a lack of security specialists with a combined understanding of technologies and security policies, highlighting a need for IAM education. Measurements for these aspects include (but are not limited to) user screening, training certificates, expected or blind field tests, training hours, budget on training, or reported incidents attributed to missing training [33, 160]. Timeliness is also a factor, as training needs to start immediately after joining an organization. Ongoing educational efforts counter its decay and keep the gained education up-to-date. In summary, education imposes initial and ongoing educational costs but eases security and operations. Education also affects the IAM software, as usability demands learnability. The target for an education metric is not the maximum. Users should be appropriately aware and trained for IAM to operate safely and securely within their responsibilities. Education is relevant for the full IAM audience spectrum.

Orphans are accounts without a valid owner. In a narrow sense, it applies to accounts without any owner. In a broader sense, however, it also applies to accounts with invalid owners, including inactive, expired, or never used ones. Orphan accounts represent violations of PoLP and compliance. Unused accounts are de facto unnecessary, and missing owners negatively affect the traceability of audit logs. Besides the number of accounts without an owner, measurements for orphan accounts may include those never logged on, inactive (e.g., after 60 or 90 days, etc.), locked or expired accounts [15, 35]. Orphan accounts are rooted in a mismatch of the identity lifecycle and deprovisioning. E.g., authentication-based protocols, like Security Assertion Markup Language (SAML) [122], OAuth 2.0 [59] or OpenID Connect [129], may provision accounts on demand and authentication but do not consider deprovisioning. More recent approaches, like System for Cross-domain Identity Management (SCIM) [69, 70, 89], separate authentication and provisioning contexts, enabling centralized real-time automatism [138], thus eliminating a technical reason for orphan accounts. However, orphans can still originate from improper identity lifecycle processes. E.g., former organization members' access is a common anti-pattern [16]. In summary, orphan accounts indicate security risks, compliance

Manuscript submitted to ACM

365 issues, and poor identity and process quality. However, a missing or ineffective deprovisioning process indicates cut
366 implementation costs at the expense of security. Because of their avoidable issues, the target value for an orphans metric
367 should be minimal. The relevant IAM audiences for an orphans metric are foremost governance and identity admins.

368 **Excess** refers to users with an intended or unintended high number of authorizations [30]. Sources for excess are
369 granting much access during joiner and mover processes or missing revokes on mover or leaver ones. Following PoLP,
370 intended excess occurs for users with administrative duties or responsibilities for sensitive data. Means for reducing
371 the resulting security risks are Privileged Access Management (PAM), education, or reorganizations. Measurements
372 for intended excess target necessary yet excessive user authorizations. Examples include the number of superusers
373 or privileged accounts and users with access to sensitive data or privileged authorizations [33, 35, 147]. Unintended
374 excess occurs on a violation of the PoLP. Because of its unintended nature, this excess negatively affects the security
375 goal. Measurements include outliers with significantly more access than their peers [35]. Orphans or the number or
376 ratio of received approvals for privileged accounts [15] also raise an unintended excess metric. Further measurement
377 dimensions for excess are the support for specifying and enforcing the PoLP [66], the duration for an (emergency)
378 deactivation of authorizations [67], or the usage for constraints avoiding excess [66]. Despite its positive effect on the
379 security goal, the target value for an excess metric is not minimal, as some excess is intentional. In this sense, careless
380 excess reduction might lead to a shortage of authorizations, hampering operations. However, unrestricted excess to on a
381 non-PoLP-compliant level harms security and compliance while staying neutral to cost reduction and operations goals.
382 From a quality perspective, a higher value for an excess metric indicates a broken identity lifecycle and low-quality
383 authorizations for identities. Finally, the relevant audiences for excess are especially governance and identity admins.

384 **Password Strength** is one of the most famous and well-studied IAM metric in literature and practice. Most users
385 notice this metric during support processes like registration or password change. Approaches for password strength
386 vary in their scaling but assign similar labels like low, medium, and high, used for visualizing the password strength
387 by password strength meters [31, 164]. The discussion of password strength effects points back to the seventies [101]
388 and maps it to operations and security goals. Hence from a theoretical perspective, a solid password strength metric
389 indicates security since passwords become harder to crack. Various measurements have been proposed [31, 151].
390 However, in practice, it is shown that humans find strong passwords more difficult to remember, which hampers
391 operations [106]. E.g., users forget their passwords and require costly help-desk or self-service processes afterward
392 [67, 160] or utilize anti-patterns like writing passwords down and leaving the notes nearby their devices. This clash of
393 security and operations goals is leading to the most recent recommendation of the National Institute of Standards and
394 Technology (NIST) to lower password strength calculations by enforcing requirements like special characters [51]. Since
395 passwords are a vital cornerstone of modern IT security, a dedicated community of the usable security realm actively
396 studies password-strength aspects, improving security and operations. For example, these studies include serious
397 games to educate about password strength [79], user perception on password expiration [53] or further strengthening
398 passwords [92, 141]. From a compliance perspective, the GDPR [37] demands "*appropriate technical and organizational*
399 *measures to ensure a level of security appropriate to the risk*" requiring state-of-the-art authentication. A securely stored
400 password is part of identity and software quality. The relevant IAM audience for password strength includes especially
401 the users themselves when they choose their password. However, the overall strength of utilized passwords is also
402 interesting for governance and operations. For governance, robust authentication methods support compliant protection
403 of unauthorized access. For operations, a strong password strength metric hints at raised support expenses.

404 **Biometric Strength** describes the security biometric authentication approaches achieve. E.g., Ryu et al. [128]
405 collected 54 performance measurements for biometric authentication schemes spanning various categories: functionality,
406 Manuscript submitted to ACM

417 security, usability, operation time, operation cost, satisfaction/acceptance, and stability. Most approaches evaluate their
418 accuracy measured with False Acceptance Rate (FAR), False Rejection Rate (FRR), and Equal Error Rate (EER) which
419 depict effects both on security and operations. The EER levels out both FAR and FRR to achieve an acceptable degree
420 of both security and operations. Ryu et al. [126] point out that other categories like efficiency, usability, security, and
421 privacy need consideration when evaluating biometrics despite the widely applied accuracy. Biometric factors can be a
422 non-intrusive control like keystroke dynamics which would not interrupt a user's task. Other biometric factors, like
423 fingerprints, interrupt a user's workflow, which hampers operations. Therefore, users perceive biometric authentication
424 schemes as more usable, supporting operations. However, security suffers from the probabilistic behavior of the error
425 rates and robustness, like for varying input device quality [111]. As for the password strength, the GDPR applies,
426 requiring a state-of-the-art authentication scheme for which biometric factors enhance as an added factor. From a
427 quality perspective, a biometric strength metric aligns with the password strength metric indicating properties for
428 high-quality identities and software. Finally, relevant IAM audiences for the biometric strength are governance for
429 compliance reasons, operations, and the users themselves.

433 **Ownership Strength** refers to authentication schemes which are based on *What You Have*. These authentication
434 approaches rely less on the chance of extracting features correctly, like biometrics. Its strength is that only an authorized
435 identity controls the transponder for authentication. This assumption is valid for devices like a Radio-Frequency
436 Identification (RFID) transponder securely attached to a key chain, in a wallet, a smartphone, or wearables always
437 nearby the identity. Some approaches tried to reach perfection in this regard and implanted RFID transponders in living
438 creatures which is practice for livestock [74]. However, privacy and security concerns make implants only feasible for
439 identification and not authentication because of cloning, destroying, or unauthorized tag reading issues, limiting their
440 usage [57, 124]. Measurements thus include, for example, lost transponders or their cloning resistance. While requiring
441 the user to authenticate is hampering operations and cost reduction, organizations raise security and comply with
442 regulations like the GDPR. Compared to other authentication factors, ownership does not rely on biometric probabilities
443 or remembering passwords while introducing the risk of losing a transponder. Relevant IAM audiences are the users
444 and operations to manage transponders and governance for overseeing compliant authentication.

448 **Multi-factor Strength** describes the IT security of combined authentication factors [111]. This approach assumes
449 that a combination of stand-alone authentication factors creates a more effective authentication scheme. Thus, rating
450 and benchmarking these factors help select relevant factors. In that sense, Ometov et al. [111] compare the factors by
451 universality, uniqueness, collectability, performance, acceptability, and spoofing. In detail, universality and uniqueness
452 indicate whether all identities can use and distinguish by the factor. Collectability denotes the ease of data collection.
453 Performance aggregates the property's accuracy, speed, and robustness. Acceptability relates to usability and gives
454 information about the practical usage of the factor. Finally, spoofing indicates the security of a factor. A multi-factor
455 strength metric aggregates these measurables. An emerging research field is adaptive authentication [4], which reacts
456 to the given authentication context, boosting operations and security goals. It works by either parametric adaption
457 or structural adaption. A parametric adaption fine-tunes parameters of the authenticator, like reducing the allowed
458 password attempts when logging in off-site. A structural adaption changes the system's structure, like requiring an
459 additional factor for the off-site login example. These adoptions require added measurements for security, usability,
460 and user context. These measurements include time, location, activity, data sensitivity, hardware measurements, user
461 preferences, etc. For compliance like the GDPR, multi-factor authentication is currently the best practice for state-of-
462 the-art authentication, recommended by the NIST [51] Authenticator Assurance Level (AAL) 2 and higher. Like other
463
464
465
466
467
468

10

Baumer et al.

469 authentication factors, interrupting a user's tasks hampers operations and raises costs. Compared to other authentication
 470 schemes, combined factors raise security. Relevant IAM audiences are users, operations, and governance.
 471

472 5.2 Policy Lifecycle Perspective

474 Policy lifecycles models are present for Role-Based Access Control (RBAC) [41, 44] and Attribute-Based Access Control
 475 (ABAC) [64]. These models include suggestions for introducing policy systems, the derivation and implementation
 476 of the policies for an organization, the policy's and system's operational maintenance or optimization [82], and final
 477 disposal. Table 2 summarizes the IAM metrics *bypass*, *adaptability*, and *constraints* for the policy lifecycle.
 478

479 Table 2. Identity and Access Management Metrics: Policy Lifecycle Perspective (own depiction based on a bottom-up derivation).
 480

Metric	Target	Cost Goal				Quality Goal				IAM Audience				
		COST	SEC	COM	OPS	PRC	ID	POL	SW	GOV	ID	POL	OPS	USR
Bypass	min	-	-	-	+	-	-	-	-	✓	✓	✓	✓	✓
Adaptability	max	+	+	+	+	+	+	+	+	✓	✓	✓	✓	✓
Constraints	target				-				+	✓	✓	✓	✓	✓

487
 488 **Bypass** describes any failure in translating a security policy to an authorization or ACP up to an effectively granted
 489 access. Any bypass thus undermines the desired IAM state and is detectable from various perspectives. E.g. for a
 490 preventive perspective, measurements include risk-based assessments of applications, cases for ignoring the actual
 491 Policy Decision Point (PDP) or Policy Enforcement Point (PEP), and the IAM system's verification capabilities as policies
 492 might grow into complex constructs leaking authorizations [66]. Alternatively, Beres et al. [15] suggest including
 493 processes and measure bypassed approval processes for provisioning. From a detective perspective, Security Information
 494 and Event Management (SIEM) systems provide relevant measurements, including the number of security incidents
 495 due to authorization management or their combinations [67], unauthorized access to confidential or sensitive data,
 496 shared credentials, multiple and quick location changes, off-hour access, etc. [97]. From a corrective perspective, quality
 497 is a relevant factor. Measurements for this corrective perspective may utilize insights from periodic Access Reviews as
 498 these indicate timeliness and accuracy of ACPs [52]. While bypasses boost short-term operations, they hamper security
 499 and compliance. A bypass also indicates poor IAM quality. Policies should authorize users appropriately. Operational
 500 support processes like self-services help users receive missing authorizations without bypasses, and IAM software
 501 should prevent bypasses. Despite its positive effect on operations, the target value for a bypass metric is minimal.
 502 Relevant IAM audiences encompass governance, identity and policy admins, and operations.
 503

504
 505 **Adaptability** is the capability of an IAM environment to reconfigure its security policies, authorizations or ACPs.
 506 Since policies reflect the status quo of an organization, changes in the environment or the organization lead to a
 507 policy decay demanding timely adaption capabilities. Adaptability is thus a prerequisite for (periodic) improvements
 508 [43, 83]. Measurements include the system's flexibility and the ease of adapting or distributing updated policies. (i)
 509 Flexibility measurements hence encompass assessments for the system's ability to cover various use cases [71, 86],
 510 features to handle policy changes, situational awareness, or capabilities to combine policies [66]. (ii) Measurements for
 511 an eased adaption may include discovery features, steps required to assign authorizations, the degree of syntactic or
 512 semantic ACPs specification support, capability and usage of delegation features, or the duration for taking adjustments
 513 [30, 66, 67]. (iii) Finally, the updated policies must also be accessible to other system components [10]. Measurements
 514 for this evaluate the distribution, source management, repositories, and retrieval of the policies [66] or the level of
 515 interoperability to communicate with other IAM systems. Connectivity to other systems via standardized protocols,
 516
 517
 518
 519
 520 Manuscript submitted to ACM

like Lightweight Directory Access Protocol (LDAP) [166], SAML [122], Service Provisioning Markup Language (SPML) [27], SCIM [69, 70, 89], or XACML [123], indicates policy adaptability. Since adaptability enables staying up-to-date with organizational needs, it benefits security, compliance, and operations goals, ultimately leading to cost reductions. From a quality perspective, high-quality IAM software provides adaptability, while policy quality benefits from eased adaptations when required. The target value for an adaptability metric is thus its maximum given that only authorized policy admins perform the updates. Finally, policy admins are the primary audience.

Constraints describe the capability of an IAM system to effectively realize restrictions for the utilized policies, enabling a safer and more secure IAM. E.g., for RBAC models, constraints restrict the number of permissions assigned to a role (cardinality) or enforce Segregation of Dutys (SoDs) [38, 133]. E.g., software engineers might be excluded from testing their written code because they may become professionally blinkered. A static SoD disallows the software engineer to test any code. A dynamic SoD prevents activating a software engineering and testing role for a given feature simultaneously. A constraints metric can thus encompass measurement dimensions regarding the number of violations or the modeling capabilities of the IAM system. Various perspectives show the number of violations. E.g., the plain total number of violations found among users [15], resulting security incidents of SoD violations [67], and discrimination for SoD violation types [35]. Furthermore, measurements for constraint modeling include the system's support for cardinalities like the PoLP [66], and SoDs like supporting policy admins to decide toxic authorization combinations [35] or providing features to specify dynamic or static SoD rules [66]. The target value for a constraint metric is ambiguous, requiring explanation. While violation measurements should minimize, modeling capabilities should meet the organization's needs. From the goals' perspective, a *good* value of the constraints metric benefits security and compliance. For the operations goal, however, an adverse effect originates from slowed down processes when SoDs require multiple parties for participation. From a quality goal perspective, modeling constraints raise the accuracy of identities and policies. The enforcement of these constraints indicates high-quality IAM software. Relevant IAM audiences are governance or identity and policy admins.

5.3 Management and Governance Perspective

While IT governance refers to looking forward and steering an organization's goals and the resulting IT resource assignments, IT management organizes and plans IT resources in the present. IT governance covers effective objectives, while IT management focuses on efficient execution. Therefore, IT governance and management are semantically close to compliance requirements [130]. Metrics to evaluate both concepts are available for IAM systems. These encompass especially *completeness*, *control*, *traceability*, *size* and *hygiene*, and are summarized in Table 3.

Table 3. Identity and Access Management Metrics: Governance Perspective (own depiction based on a bottom-up derivation).

Metric	Target	Cost Goal				Quality Goal				IAM Audience				
		COST	SEC	COM	OPS	PRC	ID	POL	SW	GOV	ID	POL	OPS	USR
Completeness	target	-		+		+			+					
Control	max	-	+	+	+									+
Traceability	max	-	+											
Size	target	-				+								
Hygiene	max	-	+		+	+	+	+						

Completeness describes the degree of actually having all requirements stated by governance decisions implemented for an organization. This metric thus indicates the effectiveness of the IAM realization. Measurements supporting completeness include the number of successful audits, audit findings [67, 147], updates, and security policies [73].

Manuscript submitted to ACM

573 Furthermore, decision-makers can use capability maturity models [138] or metric collections (like this work) to assess
574 the current status quo. The target for completeness is maximizing while considering economic properties. Completeness
575 benefits compliance goals by ensuring their realization in processes and software, which implies complete support for
576 all organizational needs and technical features. However, it also harms cost reduction as completing all requirements is
577 an ongoing investment for organizations [67]. The relevant IAM audience is primarily governance.

578
579 **Control** evaluates the overall capability to operate an IAM system within an organization. It closely links to the
580 completeness metric and compliance regulations. An IAM audience can interpret control as actually controlling an IAM
581 or just its capability. Actual control is traceable by risk and criticality assessments [68], while compliance violations
582 or education show the control capability. Measurements for actual control thus include open audit or SoD findings
583 and violations [55, 67, 84, 147] or time-spans to solve IAM related incidents as an (emergency) ACP deactivation
584 [67, 73]. From the control capability perspective, measurements include the amount of IAM entities in compliance
585 [33, 86], the amount and frequency of taken audits, updates, backups, or logs [73, 147], assessments for proper security
586 mechanisms supporting security policies [73], or evaluations of risks and criticality for IAM entities [67, 68, 115, 147].
587 The target for a control metric is maximum while not imposing micromanagement. Controlling IAM affects cost
588 reduction negatively, while security, compliance, operations, and quality benefit. The relevant IAM audiences for a
589 control metric are governance and operations for evaluating the status quo and reinforcements.

590
591 **Traceability** captures the capability for tamper-proof tracking of authorization and authentication decisions. This
592 capability is rooted in compliance regulations, like SOX [149]. However, it is also a relevant capability to understand
593 the status quo, like during debugging of ACPs or security incidents. Attackers might cover their tracks when causing
594 a security incident, raising the need for tamper-proof tracking [107]. Measurements include the presence, content,
595 frequency, and accessibility of audit logging or backups [56, 66, 73]. Also, the ratio of monitored logs, the amount of
596 documented unauthorized access attempts, or the ratio of administrative logs are named in literature [56]. From a
597 trust perspective [102], missing traceability hints at lost control. Especially compliance regulations drive traceability
598 benefiting the compliance goal to raise security regarding otherwise undetected threats [55]. Traceability thus indicates
599 software quality. However, the metric negatively affects the cost reduction goal as software for persisting and managing
600 logs must be developed, bought, and operated. Therefore, the target value for a traceability metric is its maximum while
601 staying cost-efficient. Relevant IAM audiences are governance and operations.

602
603 **Size** is a relevant metric for the management realm. It indicates the magnitude, feasibility, or requirements of the
604 challenges faced while operating an IAM system. E.g., an organization with hundreds of members requires different
605 approaches for IAM as one with tens of thousands. Thus, size measurements include managed entities and their
606 type. E.g., the size of the (IAM) repository is a basis for risk calculations [115]. Elimity [35] cover a broad range of
607 size measurements for different RBAC entities and their types, like users (newly created, disabled, or IT users) [147],
608 permissions (total amount or relative to users and roles), roles (organizational, IT, business, technical, etc.), or connected
609 applications. From a quality perspective, a broad consensus on avoiding a role explosion during role mining is present
610 [84], thus limiting the number of roles. Furthermore, role coverage indicates the number of relationships between IAM
611 entities as a measurement for a size metric [35, 85]. A bigger size metric indicates raised costs, more risk, and hampered
612 operations, negatively affecting their respective IAM goals. Size, however, also forces processes and software towards
613 more efficiency and quality to work on higher scales. The target value for a size metric is hence ambiguous, it just
614 indicates the magnitude of challenges ahead, but this does not mean it should primarily grow or shrink. The relevant
615 IAM audiences are governance, identity or policy admins, and operations.

616
617
618
619
620
621
622
623
624 Manuscript submitted to ACM

Hygiene depicts the amount of futile entities within an IAM system. These entities make operating an IAM system more cumbersome while adding no value. Measurements for hygiene thus encompass inactive, empty, or unnecessary IAM entities [15, 35, 147], error rates for processes producing unnecessary data artifacts [67], or undecided changes to IAM entities [35]. Inactive or empty IAM entities thus include orphans like described by the orphans metric before, users without access to any application, roles with no permissions or users, users with permissions assignable by roles, test entities not deleted after the test, etc. For misaligned processes, error rates on creating IAM entities measure the production of unnecessary entities. Finally, measurements for undecided changes also give insights on outdated and unclear IAM data as hanging reviews for entities impose decay of data quality and raised risks. Hygiene especially benefits identity and policy quality goals for avoiding unnecessary IAM entities. Hygiene also benefits security since decision-makers and admins get a clearer picture of the data managed by the IAM system. Operations and process goals are also positively affected as less unnecessary IAM entities improve the efficiency of technology-supported business processes. However, the cost reduction goal is affected negatively as cleanup processes or software for IAM systems require an appropriate setup and ongoing efforts. The target value for a hygiene metric is cost-efficient maximization. Relevant IAM audiences are (identity and policy) admins and operations.

5.4 Identity Quality Perspective

The quality of identities is crucial because it affects not only the results of IAM-related processes and data consumers (such as policy admins or department heads) but also the daily work routines of humans or machines, represented by digital identities in information systems. Pfizmann and Hansen [116] define identities as a subset of attributes of an individual which at least identifies this individual among any set of individuals. This well-established definition comes with two important implications: First, a person does not have one but many identities and can choose to use one identity or another for identification in a particular context. Second, a digital identity is only a set of entity attributes in an information system. The quality of digital identities is hence an instance of the data quality concept. Existing IAM standards and technologies commonly consider two kinds of semantic entities that fit the definition of digital identities: (i) Identity (or user, subject, employee, etc. [85]) data sets are a digital representation of a person. (ii) Accounts represent a person or machine in an application context with an optional association to an identity. Following these definitions and with established data quality theory, we define the quality of digital identities as the *fitness for use* of an identity or account data set in an IAM context. Table 4 summarizes the identity quality metrics presented in the following.

Table 4. Identity and Access Management Metrics: Identity Quality Perspective (derived from Cichy and Rass [24]).

Metric	Target	Cost Goal				Quality Goal				IAM Audience				
		COST	SEC	COM	OPS	PRC	ID	POL	SW	GOV	ID	POL	OPS	USR
Accuracy	max	+	+	+	+	+	+	+	+					
Timeliness	max	-	+		+									
Accessibility	target		-	-	+									
Understandability	max				+				+					
Privacy	target		+	+	-									

Accuracy means that identity data accurately represents a real-world state [154]. Common examples of inaccurate identity attributes are misspelled names or outdated addresses. Accuracy is an *intrinsic* quality dimension, so it applies to stored data, not requiring a representation to data consumers or an application context to be valid. Established data quality research considers accuracy one of the most significant data quality dimensions. Maintaining accuracy is a particular challenge [145] since organizations usually utilize a heterogeneous identity management infrastructure (e.g.,

Manuscript submitted to ACM

677 HR systems, directories, application systems, or staging databases [42]). Accuracy affects the quality of work results
678 for all data consumers, leading to far-reaching repercussions of poor identity accuracy: Deficiencies in identity data
679 can cause incorrect authentication or authorization decisions, inoperable accounts or application systems, or process
680 interruptions (e.g. when an approver needs to double check implausible data before proceeding). As a result, we define
681 the target function of this metric as maximization and consider all IAM goals as positively affected by high accuracy.
682

683 **Timeliness** determines whether data is up-to-date, either in terms of accuracy or being able to serve a task at hand
684 [155]. While timeliness for accuracy is primarily relevant for preventing identity data errors [135, 144], timeliness for
685 serving a task at hand is essential if data loses validity after some time (e.g., real-time status data). E.g., outdated identity
686 data is a problem for engineering ACPs, and for ABAC also for correct authorization decisions in real-time [85]. In
687 these cases, the timeliness of identity data greatly benefits policy quality, security, and operations goals. However,
688 keeping identities up-to-date over a more extended period requires regular effort, contrasting the cost reduction goal.
689 The primary audience for this metric is identity admins for keeping identities up-to-date. Policy engineers relying on
690 this data and users providing the data input are also relevant audiences.
691

692 **Accessibility** issues can prevent data consumers from working with data for several reasons [145]: IAM infras-
693 tructures typically comprise a variety of (identity) data sources with varying data schemas and integration levels [42].
694 Although this data may be technically accessible, its distribution and lack of standardization are common reasons
695 for data consumers to perceive it as inaccessible or unprocessable. Poor data understandability and inconsistent data
696 representation can have the same effect. Another typical availability issue is incomplete data: Supplementary identity
697 information such as access logs can often only be collected to a limited extent or only for parts of the identity manage-
698 ment infrastructure [83]. Similarly, large data volumes (e.g., user activity logs) can easily lead to timeliness problems
699 (e.g., when collected in batches over longer intervals). As a result, the computed data is unavailable when needed. At
700 last, authorization constraints impair accessibility intentionally: Both confidentiality and privacy are essential security
701 goals that make it necessary to restrict access to identity data. The goal is thus not to maximize data accessibility but to
702 manage it accurately based on an organization's security policy [134]. High accessibility positively affects operations
703 but harms security and compliance goals. Inaccessible identity data can prevent software from functioning correctly
704 (e.g. when relevant steering attributes are unavailable to authentication or authorization mechanisms). The primary
705 audiences for this metric are identity admins and governance.
706

707 **Understandability** determines whether data consumers can understand and use identity data [119, 145]. This quality
708 property is critical when human data consumers need to process unfamiliar digital identities (e.g., policy admins or
709 deciders in a self-service process). Sufficient understandability is a common requirement when IAM processes integrate
710 users throughout an organization (e.g. when applying domain knowledge in approval processes). Examples of improving
711 identity understandability are speaking display names rather than cryptic identifiers, pseudonyms, or proper description
712 texts in languages the data consumers understand. Understandability is closely related to data accessibility and can
713 conflict with privacy (e.g., for anonymized or pseudonymized identity processing requirements). Altogether, identity
714 understandability benefits operations and policy quality. The primary audience includes identity admins, ensuring that
715 humans unfamiliar with the identity data can comprehend them.
716

717 **Privacy** is a communication restriction of personal identity data regarding the extent, processing, and moment [158].
718 In the context of IAM, privacy is relevant for data management and software (see Section 5.7). When working with
719 digital identities, data privacy requirements of users or regulators require consideration [54, 55, 71, 73, 86, 135, 144]. In
720 addition, a large number of standards and frameworks define corresponding principles. E.g., the GDPR [37] requires user
721 consent and minimal processing of personal data. Furthermore, personal data requires special security measures that
722
723
724
725
726
727
728

prevent unauthorized reading or leaks. Wagner and Eckhoff [153] present 79 privacy-related metrics or measurements regarding uncertainty, data similarity, indistinguishability, time, error, accuracy, adversary's success probability, and information gain or loss. Further typical privacy metrics are anonymity and linkability [25, 26]. Identity privacy is vital for legal compliance, and identity data minimality also benefits security, as only accessible data can leak during a security incident. At the same time, missing or inaccessible data hampers processing for productive tasks and negatively affects operations [145]. The goal of IAM is thus not privacy maximization but ensuring *sufficient* privacy under productive operation. Privacy metrics are essential for governance and trusting users for careful personal data management.

5.5 IAM Process Quality Perspective

Processes are crucial for IAM, reflected by the broad overview of process frameworks given by literature: (i) Identity lifecycle models focus on joining, moving, and leaving identities or the provisioning and deprovisioning in distributed application landscapes [75, 112, 125, 159]. (ii) Policy lifecycle models focus on creating, enforcing, maintaining, and deleting policies, their provisioning in a decentralized application landscape, managing and enforcing responsibilities, and requesting missing authorizations as self-service [44, 64, 82, 90]. (iii) Finally, IAM capability maturity models take a broader perspective and include processes such as help desks, approvals, recertifications, and logging and tracking of activities within an IAM infrastructure [112, 138, 159]. Process quality is no original IAM concept but is well-studied by business and information systems research. We thus apply the widely acknowledged Quality Evaluation Framework (QEF) [60] for IAM process quality. QEF defines a quantitative assessment of process quality based on six essential quality dimensions: Performance, efficiency, reliability, recoverability, permissibility, and availability. Table 5 summarizes the metric alignments for this perspective. For simplicity, we use the joiner process as a running example.

Table 5. Identity and Access Management Metrics: Process Quality Perspective (derived from Heidari and Loucopoulos [60]).

Metric	Target	Cost Goal				Quality Goal				IAM Audience							
		COST	SEC	COM	OPS	PRC	ID	POL	SW	GOV	ID	POL	OPS	USR			
Performance	max	+			+	+			+								
Efficiency	max	+			+	+			+								
Reliability	max	+		+	+	+			+								+
Recoverability	max	+			+	+			+								
Permissibility	max			+	+				+								
Availability	max	+			+	+			+								+

Performance determines the efficacy of a process. An example is the process of an identity joining an organization. The new identity should be ready to work as soon as possible. A manager thus needs to order required authorizations for the identity from the IAM team (e.g., via an IAM self-service [63]). This activity involves a manual review of the requested authorizations by the IAM team. (i) The first performance criterion is *throughput* as the number of processed elements in a given period. The example applies throughput to the number of authorization requests the IAM team can review daily. (ii) *Cycle time* refers to the idle time during the preparation and execution of a process. Since cycle time results from process delays, a minimum is desirable. For example, the IAM team may take time to start the authorization check. Added cycle time might occur during the request review. For example, when the IAM team needs to wait for an appointment with an application admin to determine the criticality of a requested permission. Beres et al. [15] provide further meaningful IAM examples for process throughput and lag. (iii) *Timeliness* refers to the time between initiating and terminating a process. While a process starts with perfect timeliness of 0, any delay will reduce it. In the example, the review due time might be specified as two days, meaning that it is sufficient for a manager to order authorizations

781 two days before the join. If the processing is started after one day and takes another two days to complete, the timeliness
782 is -1 day as the processing delay is one day. (iv) Finally, *cost* determines the amount of money required for the process
783 execution. The example includes the hourly personnel costs required for the IAM team to review authorization requests.
784 Operations and cost reduction are positively affected by maximizing process performance. Relevant audiences for
785 process performance are operations and governance, ensuring smooth and cost-efficient processes.
786

787 **Efficiency** relates process performance to its economic viability. (i) *Time efficiency* indicates the ratio of the lead time
788 of a process to the expected execution time. It is thus related to timeliness, but timeliness is an absolute value designed
789 to meet a hard time limit. In contrast, time efficiency looks at the relationship between planned and actual time. For the
790 running example, the expected duration of an activity review is two days, but the review takes three days to complete.
791 According to the QEF, the time efficiency of the review activity is $2/3 = 66\%$. While timeliness tries to prevent negative
792 outliers altogether, this relative depiction is more helpful for minimizing process runtime. (ii) Similarly, *cost efficiency*
793 is defined as the ratio of planned costs to required costs for an activity execution. Like process performance, maximized
794 process efficiency is desirable as it positively affects operations and cost reduction. For the same reasons, the audiences
795 for process efficiency are governance and operations, ensuring smooth and cost-efficient processes.
796

797 **Reliability** concerns the number of failures occurring during process execution in a given period. For example, a
798 created but not activated account of a joining identity may prevent subsequent logins. (i) *Reliability* describes the
799 probability of failure in a given period. For example, this probability may describe the likelihood of all joining identities
800 having their accounts properly provisioned last month. (ii) *Failure frequency*, in contrast, describes the number of
801 failures in a given period: e.g., the number of joining identities that could not log in last month. High process reliability
802 enables operations and prevents avoidable rectification costs, security risks, and compliance problems (like excessive
803 authorizations). Since failures in lifecycle processes are likely to cause errors for entities, their reliability positively
804 affects identity and policy quality. The key audiences for process reliability metrics are governance and process owners.
805

806 **Recoverability** considers spent failure recovery time. E.g., an admin might spend an hour analyzing reasons for the
807 failed logins and activating an inactive account. The QEF recoverability definition includes two intervals: First, *time to*
808 *failure* defines as the time between two failure recoveries (e.g., time since the last failure). Second, *time to recovery*
809 is the time for failure rectification. Building on these variables, *maturity* defines the relative time spent without a failure:
810 E.g., the time for process execution without rectifying errors. Thus, a high maturity leads to recoverability, meaning
811 that processes execute with minimal expenses on failure recovery. High recoverability positively affects operations and
812 cost reduction. Like reliability, the audiences for recoverability are governance and process owners.
813

814 **Permissibility** determines whether the execution of a process conforms to defined authorizations or violates them.
815 QEF differentiates between two kinds of authority: (i) *Input authority* determines whether an activity of a process that
816 accesses a process input (e.g., a physical resource or data) is authorized to do so. For example, reviewing authorization
817 orders for a joining identity might require the decider to investigate the joiner's responsibilities. The decider might
818 have authorizations to inspect the joiner's department and job title but not the authorizations of the colleagues, which
819 is helpful context information [78]. (ii) *Activity authority* determines if an actor that executes an activity is authorized.
820 For example, a responsible decider might not have the security clearance to make authorization decisions for other
821 identities. From the perspective of an individual process, the process permissibility should be 100%, so the process
822 accesses all required resources by responsible actors without violating authorizations. Process permissibility is strongly
823 linked to ACP authorization accuracy. It is relevant from a security perspective [15] and thus positively affects security
824 and compliance goals. The audience for permissibility includes governance and policy admins.
825

826
827
828
829
830
831
832 Manuscript submitted to ACM

Availability describes the time a process executes as intended. Unlike recoverability, which considers outages caused by failures in the process execution, availability is concerned with outages caused by the unavailability of required inputs. Such input can be physical resources and information required for the process. The *time to shortage* of a process is the time that goes by from one shortage to another (e.g., the time processes execute as intended). The *time to access* is the time required until shortage remediation. Consequently, QEF defines *availableness* as the ratio of time to a shortage with the total process execution time (time to shortage + time to access): E.g., the ratio of time for a process without shortage. For example, consider that a self-service [63] is not available 10% of the time when a manager orders authorizations. The availableness is thus 90%. High availability is crucial for operations, while low availability is a cost driver. Availability is especially relevant for governance and operations, contributing to interruption-free execution.

5.6 IAM Policy Quality Perspective

Various policy types are present in organizations, such as regulatory or internal requirements. For the IAM domain, policies commonly refer to Access Control Policies (ACPs). ACPs are machine-processable rules, determining allowed user access [131]. Access Control Models (ACMs), such as RBAC [114, 133] or ABAC [64, 143] define the ACP structure. Consequently, ACP correctness is a prerequisite for ensuring correct policy mappings of the access control mechanism: While loosely defined policies are security risks, too restrictive ones hamper operations by preventing user tasks. ACP quality can be considered a quality concept in its own right. Nevertheless, approaches exist in scientific research to place it in a larger frame of reference as an instance of data [83] or software quality [161, 162]. Table 6 summarizes metric alignments of this section.

Table 6. Identity and Access Management Metrics: Policy Quality Perspective (derived from Kern et al. [83]).

Metric	Target	Cost Goal				Quality Goal				IAM Audience				
		COST	SEC	COM	OPS	PRC	ID	POL	SW	GOV	ID	POL	OPS	USR
Authorization Accuracy	max		+	+	+	+				✓				
Evaluation Runtime	min				-	-								✓
Maintainability	max	+								✓				
Understandability	max	+		+		+				✓				
Complexity	min	-												
Redundancy	min	-	-		-	-								
Automation	max	+												
Policy Attribute Quality	max					+		+						

Authorization Accuracy describes the precision of policies depicting authorizations defined by organizational security policies [12]. Two kinds of errors cause a deviation from correct authorizations: Excessive authorizations and missing authorizations. Excessive authorizations describe more granted permissions than necessary, leading to security risks. Missing authorizations describe a shortage of needed authorizations, harming operations. The difficulty of measuring authorization accuracy is determining authorizations a subject *should* have since this requires a profound understanding of a user’s tasks, the implications of an authorization, and the requirements of internal security policies. As a result, determining inaccurate permissions often requires context knowledge from human domain experts. Nevertheless, the permissibility of processes reflects policy accuracy (e.g., approval accuracy [15]) as it positively affects operations and security. Regulatory compliance also requires sufficient accuracy of authorizations. Since missing authorizations cause interruptions, authorization accuracy is positively related to process quality. Policy accuracy also influences future policy quality, as maintenance likely produces inaccuracies (e.g., when using bottom-up methods such as mining or learning-based automation [83]). The audience includes governance and policy admins.

885 **Evaluation Runtime** determines the evaluation speed for making an authorization decision based on the present
886 policies [83]. It is correlated to the response time of the underlying access control system [66]. This quality dimension
887 is essential for XACML, which defines a reference architecture, including a real-time policy evaluation. A centralized
888 access control mechanism like XACML [123] is an operations bottleneck. A low evaluation runtime thus improves
889 operations, while a high one hampers it. The audience for this metric includes policy admins and operations staff.

891 **Maintainability** determines the ease of keeping present policies up-to-date. It indicates the ability to maintain a
892 high policy quality long-term and the administration ease for an access control system [66, 82]. ACP maintainability
893 hence affects the adaptability of an IAM system positively (see the Adaptability metric of Section 5.2). While IAM
894 policy literature often uses the terms maintainability and understandability interchangeably, many factors influence
895 the policies' maintainability, including their accuracy, human comprehensibility, and error-proneness. Since these
896 factors are subject to quantification, policy maintainability aggregates the following maintainability-related metrics.
897 Proper maintainability is a prerequisite for cost-efficient and effective policy maintenance, correlating positively to cost
898 reduction and policy quality. The audience includes governance and policy admins.

901 **Understandability** of ACPs is crucial for their management and maintenance. Humans need to understand the
902 meaning and implications of policies to determine which policies should affect certain subjects or objects. Examples
903 are decisions for self-service requests [63] or the correctness of an existing policy during access reviews [52]. Access
904 control systems often offer verification and compliance support functions [66] or visualization techniques [93] to
905 improve ACP understandability. A commonly named requirement for well-understandable policies is their semantic
906 meaningfulness [30]. RBAC directly tackles semantic meaningfulness as it understands roles as semantic constructs
907 representing a real-world concept, such as a departmental affiliation or a security clearance. However, the semantic
908 meaningfulness of policies does not depend on the underlying ACM since both RBAC and ABAC allow the modeling of
909 policies with significant or little semantic meaning. While the semantic meaning is difficult to grasp algorithmically, Xu
910 [163] proposes to indicate the degree of semantic meaningfulness by measuring their accordance with semantically
911 meaningful attributes. Other factors helpful for coherent policies are low complexity, understandable attributes (such
912 as meaningful names or descriptions), and the absence of errors in existing policies. Policy understandability is vital
913 for policy maintainability and its impact on IAM processes, like authorization self-services or access reviews. It hence
914 contributes to cost reduction and future policy quality. Furthermore, human-understandable policies are also crucial for
915 compliance since an auditor needs to approve the correctness and proper policy management by selecting random
916 policies for inspection. The primary audience for this metric includes governance and policy admins.

921 **Complexity** determines the number of data elements a policy includes. For example, numerous role mining
922 approaches measure complexity with Weighted Structural Complexity (WSC) [99, 163], a prominent approach in IAM
923 literature. A low complexity helps humans to understand policies and makes them easier to process algorithmically.
924 Complexity is a popular metric for maintainability because of its straightforward evaluation [30]. Thus role mining
925 algorithms use complexity for creating *good* policies [84]. Due to its maintenance-related function, low policy complexity
926 is mainly relevant for reducing costs and enabling high future policy quality. Its audience includes policy admins.

928 **Redundancy** occurs if policies define an authorization more than once. E.g., a subject receives authorizations via
929 two roles, or an ABAC policy includes three concurrent rules forbidding a specific user authorization. Redundancy is
930 problematic because it is a common error source: If a policy set is updated, redundant definitions can prevent a single
931 change from becoming effective unless the update considers redundancy. In the worst case, redundant authorizations
932 prevent authorization revocations in emergency revocation cases. Policy redundancies contradict security and IAM
933 process quality goals. Suppose a policy update error occurs because of redundant negative authorizations. In that
934
935
936

Manuscript submitted to ACM

case, an affected user will not receive the required authorizations, harming the operations goal. Besides their potential for error, redundant definitions cause unnecessary complexity and confusion during policy maintenance, negatively affecting cost reduction and policy quality goals. The audience includes policy admins.

Automation describes the grade a policy adapts to a changing environment without manual updates. While RBAC defines static roles, many automation mechanisms automatically update roles when real-world changes occur [83]. Since ABAC policies utilize attributes, environment changes do not affect them if the changes propagate toward the attributes. Thus, policy automation is crucial for operational efficiency and helps keep policies updated longer. However, automation cannot make maintenance obsolete since every automation mechanism operates in a limited scope and cannot cope with unexpected environmental changes. Moreover, an automation mechanism is subject to quality considerations like the original policy and can contain errors or be outdated over time. Policy automation thus benefits cost reduction and policy quality due to its maintenance support. The audience includes policy admins.

Policy Attribute Quality does not directly influence the decision of the policy evaluation mechanism. However, their existence and accuracy are vital for policy maintainability: First, human-understandable attributes (like a descriptive display name or an extended description) are decisive for the human understandability of policies [35]. Second, attributes that define responsibilities (like a policy owner) enable humans to determine a point of contact for questions like determining the policy's correctness. Third, sufficient policy attribute quality is a prerequisite for policy-related processes (like access reviews), requiring a responsible human decider [85]. As a result, it impacts the quality of IAM processes and the future policy quality. A policy attribute quality metric is primarily relevant for policy admins.

5.7 IAM Software Quality Perspective

IAM technologies offer services supporting IAM measures or processes. They comprise standards and specifications (e.g., SSO standards like Kerberos [105] and SAML [122] or data exchange standards like SCIM [10, 69, 70, 89], SPML [27], or XACML [123]), software solutions (e.g., directory services), and cyber-physical systems (e.g., biometrics). Mladenova [98] shows that software quality is well-studied with a consensus on metrics [19, 34, 50], also addressed by ISO 25010 [76]. This work connects the ISO 25010 metrics with IAM. Table 7 summarizes the metric alignments for this section.

Table 7. Identity and Access Management Metrics: Software Quality Perspective (derived from ISO 25010 [76]).

Metric	Target	Cost Goal				Quality Goal				IAM Audience				
		COST	SEC	COM	OPS	PRC	ID	POL	SW	GOV	ID	POL	OPS	USR
Functionality	target	-	+	+	+	+	+	+	+	✓				✓
Performance	max	-			+	+	+	+	+	✓				✓
Compatibility	target	-	+	+	+	+	+	+	+	✓				✓
Usability	max	-	+	+	+	+	+	+	+	✓				✓
Reliability	max	-	+		+	+	+	+	+	✓				✓
Security	max	-	+	+	-	+	+	+	+	✓				✓
Maintainability	max	-	+		+	+	+	+	+	✓	✓		✓	✓
Portability	max	-			+	+	+	+	+	✓				✓

Functionality describes whether an IAM system comprehensively, correctly, and appropriately provides its functionalities. For IAM systems, these functionalities need to back governance, identity or policy lifecycles, and support processes. E.g., Hu and Scarfone [66] dedicate a study on 27 comprehensive and appropriate Access Control (AC) system functionalities along measurements. Examples include auditing, eased ACPs discovery and assignment, integration flexibility, connecting and adapting policies of the IAM software into present systems, granularity and scopes of control, avoidance of bypasses or SoDs violations, or supportive dimensions like connectivity (imports and exports),

Manuscript submitted to ACM

989 compatibility, verification functionalities, user interfaces, and Application Programming Interfaces (APIs), etc. For the
990 identity lifecycle, the presence of functionalities for the joiner, mover, and leaver processes alongside support processes
991 like authentication (knowledge-, biometric-, ownership, or multi-factor-based [109]), help desks or self-services are a
992 valid basis for measuring the functionality of an IAM system [42, 112]. Furthermore, incorrectly implemented IAM
993 functionalities might lead to trust issues [102]. Following this concept, bug trackers or the Common Vulnerability
994 Scoring System (CVSS) [39] also estimate the functional correctness of IAM software. An additional consideration for a
995 functionality metric is organization-specific comprehensiveness. E.g., an organization that does not (need to) utilize
996 SoD policies also does not require their functionality. For this organization, a functionality metric thus should not
997 consider SoD. From an IAM goal perspective, a comprehensive, correct, and appropriate set of IAM functionalities acts
998 as a cost factor. However, these functionalities benefit all other IAM goals since powerful tools raise security, enhance
999 operations, and enable high-quality controls for processes, identities, policies, and other software. The target value
1000 for a functionality metric thus should be at maximum (for a consideration of a comprehensive functionality set) or a
1001 cost-efficient target (for required functionalities of an organization). The relevant IAM audiences for a functionality
1002 metric are governance and operations for providing IAM functionalities both on a management and operational level.

1006 **Performance** represents an efficient usage of system resources. As a metric, performance distinguishes between time
1007 behavior, resource utilization, and capacity. (i) Time behavior considers response and processing times or throughput
1008 rates. (ii) Resource utilization measures the amount and types of used resources. (iii) Capacity assesses whether software
1009 limits meet the organizational requirements [76]. IAM literature especially discusses performance for Federated Identity
1010 Management (FIM). Thus, response times [66, 135, 167], resource usage [135], amounts of requests [115], available
1011 bandwidth [144], connected systems [147], or scalability [71, 86, 135] are examples for measurements across literature.
1012 A high value for a performance metric negatively affects the cost reduction goal in the short term because proper
1013 performance upkeeping requires ongoing assessments and adjustments of used software. However, other goals benefit
1014 from performing software, like operations or process and software quality. A performance metric thus targets a
1015 maximization strategy. The relevant IAM audience for performance are governance and operations officers because of
1016 their responsibility for selecting and running the software components.

1020 **Compatibility** as a metric assesses the capability of a system or software component to exchange information
1021 and perform functions in a shared environment. Compatibility thus encompasses co-existence and interoperability.
1022 Co-existence describes the ability to execute required functions efficiently without harming other components in a
1023 shared environment. Interoperability describes the capability to provide, exchange, and use the information of the
1024 environment components [76]. For IAM approaches, trust metrics include co-existence by assessing the trustworthiness
1025 of a user's identity across an FIM covering a wide range of aspects [3, 32, 49, 102, 103, 118, 132]. Measurements for
1026 a healthy co-existence may include ex-post assessments, like unauthorized access to confidential data as Personally
1027 Identifiable Information (PII) or financial details [97]. IAM interoperability utilizes standardized protocols but also
1028 covers organizational factors like a seamless software integration (e.g., SSO) [71]. Interoperability measurements thus
1029 encompass the number of connected systems [147] or the exchange protocol support [66, 86, 142], like LDAP [166],
1030 SAML [122], SPML [27], SCIM [69, 70, 89], or XACML [123]. From the IAM goals perspective, compatibility undermines
1031 the cost-reduction goal in the short term since its elevated development expenses. Other IAM goals benefit from
1032 compatibility as co-existence considers security, while interoperability supports operations and identity or policy
1033 quality. Overall compatibility positively affects IAM software. A compatibility metric has an ambiguous target value,
1034 as improving compatibility beyond required organizational use cases is ineffective. The relevant IAM audiences for a
1035 compatibility metric are governance and operations since these select and operate compatible software.

1036
1037
1038
1039
1040 Manuscript submitted to ACM

1041 **Usability** of IAM systems defines as being an appropriate tool for IAM tasks. Users thus process their tasks
1042 efficiently and effectively while staying satisfied working with the system [21]. Moreover, the ISO25010 [76] details
1043 other properties, like appropriateness, recognizability, learnability, operability, user error protection, user interface
1044 aesthetics, and accessibility. Literature addresses these properties, especially for authentication as an IT security focal
1045 point. Authentication measurements include task efficiency (e.g., time to authenticate), task effectiveness (e.g., login
1046 attempts), and user preferences [81, 111]. For example, task efficiency reflects by the processing duration of user requests
1047 and task effectiveness by the number of failed authentication requests [67]. In that sense, Staite and Bahsoon [144]
1048 point out the role of usability by an understandable and common authentication interface (e.g., SSO [66]). Additionally,
1049 influences based on age [106], cognitive abilities [13, 14] or device characteristics [137, 152] are also studied, which
1050 need consideration as measurements for a usability metric [95]. In this context, the Fail To Enrol (FTE) rate highlights
1051 that a one-size-fits-all authentication factor does not exist. E.g., a person with deformed eyes might have trouble using
1052 an iris scan, while manual labor might wear down a fingerprint [5, 100], which can render particular authentication
1053 schemes unusable [91]. Despite these specific usability measurements for authentication, more general usability metrics
1054 like System Usability Scale (SUS) [21, 22] and Quality in Use Integrated Measurement (QUIM) [139, 140] have shown
1055 fitness to express IAM usability [20, 127]. SUS includes ten well-studied statements with a Likert scale and quicky to
1056 set up [6, 88]. Alternatively, QUIM proposes a system of ten usability factors: efficiency, effectiveness, productivity,
1057 satisfaction, learnability, safety, trustfulness, accessibility, universality, and usefulness, again supported by 26 potential
1058 measurements. Furthermore and besides authentication, usability for IAM systems is also discussed for administrative
1059 purposes. While general usability metrics, like SUS and QUIM, also apply, specific measurements are discussed in the
1060 literature, again spanning effectiveness, efficiency, and satisfaction. For example, measurements include the effective
1061 outcome of the modeled policies and affiliated verification functionalities [66], efficiency assessments for consistent
1062 user interfaces and experience [66, 86], and satisfaction or the reduction of administrative efforts in general [67]. While
1063 introducing usability implies initial costs, it benefits all other goals because IAM tasks are executed more effectively,
1064 efficiently, and with higher user satisfaction. Since the GDPR [37] also demands appropriate state-of-the-art security
1065 tools, a baseline for usability is required. The metric thus targets its maximum. The relevant IAM audience of a usability
1066 metric includes especially governance and operations for selecting and operating (usable) tools.

1073 **Reliability** describes whether software components perform specified functions in specified conditions. Furthermore,
1074 a reliability metric includes maturity, availability, fault tolerance, and recoverability. (i) Maturity denotes whether a
1075 software component fulfills requirements under usual operation. (ii) Availability describes if a software component is
1076 accessible and operational for its required usage. (iii) Fault tolerance assesses the degree to which a software component
1077 operates as intended, although hardware and software faults are present. (iv) Finally, recoverability measures whether a
1078 system returns to its desired state after an interrupting event [76]. For IAM, maturity measurements encompass maturity
1079 models allowing a comprehensive overview of a system's maturity [138]. For FIM, maturity assessments also include
1080 connected applications [147] or the support of standardized exchange protocols [66, 142]. Examples of availability
1081 measurements are the amount of time (or its percentage) a system is available for regular operation [73], its scaling
1082 capability, or potential points of failure [86]. In addition, IAM availability also covers exchanging (authentication) data,
1083 as measurements for authentication and of the identity (profile) point out [144]. For fault tolerance, wrong access
1084 control decisions due to insufficient data or the service's access control independence assess the system's behavior
1085 in tense situations [71, 135]. Besides, an assessment for potential weaknesses leading to these failures is also relevant
1086 [73]. Another aspect of fault tolerance is trust [135]. Rigorous research on trust metric in the context of FIM proposes a
1087 wide range of aspects [3, 32, 49, 102, 103, 118, 132]. Finally, recoverability measurements include the time a software
1088
1089
1090
1091
1092

Manuscript submitted to ACM

1093 component takes to recover from an interruption or the speed for rolling out a backup [73]. In this sense, measuring
1094 the frequency of backups for specific data types, such as critical or audit trail data, is also reasonable [73]. Reliability
1095 positively affects IAM goals by reducing the risks of blackouts and fostering operations for users requiring reliable
1096 services. Users also expect reliability for processes and software. The target for a reliability metric is thus its maximum.
1097 The IAM audience includes governance and operations for choosing and operating reliable software.
1098

1099 **Security** assesses a system's protection level appropriate to security policies and authorizations. The ISO25010
1100 [76] details security with confidentiality, integrity, non-repudiation, accountability, and authenticity. Please note that
1101 the Confidentiality, Integrity, Availability (CIA)-Triad, another common security model, also includes availability,
1102 already covered within reliability. (i) Confidentiality describes that data is accessible only for authorized identities.
1103 (ii) Integrity assesses preventing unauthorized modification of programs or data. (iii) Non-repudiation describes the
1104 execution of an action or event, such as it is indisputable afterward. (iv) Accountability measures the degree of
1105 uniquely traceable actions of its responsible identity. (v) Authenticity denotes whether an identity is proven the one
1106 claimed [76]. For IAM, confidentiality measurements cover privacy [1, 153, 157] affecting identity attribute disclosure
1107 [54, 55, 71, 73, 86, 135, 144]. A comprehensive survey [153] highlights the depth of this research area by categorizing 79
1108 privacy metrics or measurements in uncertainty, data similarity, indistinguishability, time, error, accuracy, adversary's
1109 success probability, and information gain or loss. Privacy-enhancing IAM internalizes these metrics [40, 58], like
1110 linkability or anonymity [25, 26]. Besides that, privacy also affects (biometric) authentication schemes since further
1111 (biological) information might leak [95]. Despite its low acceptance, an illustrative example is utilizing body odor as a
1112 biometric factor, revealing further personal information, like recent activities or possible diseases [17, 48]. Additional
1113 IAM confidentiality measurements are the number of identity data attacks, like eavesdropping, skimming, or snooping
1114 attacks [55]. For integrity, wrong access or authorization decisions based on insufficient or altered data or software
1115 indicate integrity [135]. Additionally, the number of integrity-related attacks, like data tampering, privilege escalation,
1116 spoofing, identity theft, or luring attacks, make valid measurements [55]. For non-repudiation and accountability,
1117 the traceability described earlier is applicable. Additionally, access reviews [52, 78] or (blockchain-based) audit trails
1118 [107] document non-reputable and accountable authorization decisions, making their frequency and scope relevant
1119 to measure. Authenticity utilizes the users' password, biometric, ownership, and multi-factor strength metrics, while
1120 services also need to authenticate themselves [86]. Attacks affecting authenticity include identity theft, replay, or
1121 spoofing attacks [55], whose occurrences are relevant measurements. Finally, general security metrics like CVSS [39]
1122 influence IAM. Efforts to provide security raise short-term costs. Additionally, high security often influences operations
1123 negatively because of interrupting security checks like authentication. However, security benefits other IAM goals,
1124 like reduced security risks caused by leaking data or unauthorized modifications or compliance requirements by
1125 using state-of-the-art technologies [37, 40]. From the quality perspective, identities and policies are concerned with
1126 confidentiality and unauthorized modifications. The security metric target is thus its maximum. Finally, the relevant
1127 IAM audiences for a security metric are governance and operations for selecting and operating secure software. Other
1128 audiences are identity and policy admins for modifying and reading identity and policy data.
1129

1130 **Maintainability** represents the ability of a system to modify, correct, or adapt to changes in environment or
1131 requirements. Other characteristics are modularity, reusability, analysability, modifiability, or testability. (i) Modularity
1132 denotes if a system comprises discrete and easily exchangeable components. (ii) Reusability describes whether a
1133 component can or could be used more than once in a system. (iii) Analysability measures a software component's
1134 features for diagnosing its impact, deficiencies, causes of failure, or parts requiring modification. (iv) Modifiability
1135 denotes the capability of a software component for modifications without causing defects or degradation. (v) Testability
1136
1137 Manuscript submitted to ACM

1145 describes whether operators can set up, perform and evaluate test criteria for the software component [76]. For IAM
1146 software, modularity is evident as multiple components need to fulfill various interdependent tasks [121]. XACML
1147 [123], the reference architecture for ABAC, highlights modularity by distinguishing crucial and replaceable software
1148 components, like the PEP, PDP, Policy Information Point (PIP), or Policy Administration Point (PAP). The horizontal
1149 scope for IAM software (e.g., supporting single or multi-host deployment), the presence of APIs, capabilities to import and
1150 export policies, and risks for vendor lock-ins are pointers for modularity measurements [66, 71]. Moreover, reusability
1151 builds upon modularity. For IAM, measurements should include standardized APIs, like LDAP [166], SAML [122],
1152 SPML [27], SCIM [69, 70, 89], or XACML [123]. Modularity and reusability intersect with compatibility and portability
1153 metrics. Analysability measurements include capabilities to trace modifications, as the traceability metric describes.
1154 Effective Graphical User Interfaces (GUIs) and verification functionalities also foster analysability [66]. Modifiability
1155 measurements include capabilities to configure policies into existing systems and to adapt them during their lifecycle
1156 [66]. This characteristic thus intersects with the adaptability metric. Additionally, modification or integration costs
1157 for software show its modifiability [71, 135]. Finally, testability measurements include verification features for policy
1158 engineers [66]. Maintainability is a driver for short-term costs. However, other IAM goals benefit from maintainability.
1159 Security, operations, and process quality goals benefit, especially by continuously adapting software components to the
1160 current organizational needs or the environment. Policy quality also profits from the testability aspect of maintainability.
1161 Overall, the target for maintainability is thus maximum. The IAM audiences are governance and operations for selecting
1162 and operating the software. For policy administrators, the testability characteristic is relevant.

1167 **Portability** describes an operator's ease of installing, adapting, and replacing software components on hardware or
1168 organization-specific environments [76]. For installing (and uninstalling), measurements include Operation System (OS)
1169 compatibility [66, 86, 104]. The ease in adapting includes metrics like the maintenance and adaptability metrics described
1170 before [66, 86]. Assessments for replacing of IAM software include the capability to provide a consistent experience
1171 across multiple software components [54, 55, 86, 104], avoidance of vendor lock-ins or usage of standardized protocols
1172 [54, 104], like LDAP [166], SAML [122], SPML [27], SCIM [69, 70, 89], or XACML [123]. Furthermore, a portability
1173 metric is especially relevant for recent approaches for cloud-based IAM. Criteria thus include using standards to ensure
1174 a consistent user experience, lightweight and scalable communication between software components, or the support of
1175 comprehensive native platform languages [104]. Portability negatively affects short-term cost reduction because of
1176 raised vendor expenses for specific OSes or environments. However, operations, process quality, and software quality
1177 goals benefit from the eased introduction of IAM software. The target value for portability thus is its maximum. The
1178 relevant IAM audiences are governance and operations for selecting and operating portable software.

1183 6 DISCUSSION

1185 The following discussion comprises a reflection on the taken perspectives for introducing the IAM metrics, further
1186 insights on the IAM goals and their interdependence, understandability, and limitations.

1188 6.1 Perspectives

1190 This work utilizes seven perspectives to present the proposed IAM metrics. The perspectives derived by IAM processes
1191 and goals are the identity lifecycle, policy lifecycle, governance, process quality, identity quality, policy quality, and
1192 technology quality. This section discusses the implications of each perspective and its interdependence.

1194 Insights on the identity lifecycle perspective encompass an area of conflict between cost-cutting and identity quality
1195 and distinctions for authentication strength. Although costs drive a smooth identity lifecycle, naive cost-cutting comes

1196

Manuscript submitted to ACM

1197 at the expense of identity quality. For example, orphans and excessive authorizations are rooted in anti-patterns,
1198 like only focusing on adding identities and authorizations but not considering termination. This cost-cutting boosts
1199 process performance and operations at the expense of security, compliance, and identity quality. Another aspect is the
1200 authentication strength, spanning four metrics within this work. While these have the same base effects on IAM goals,
1201 they differ in their details. E.g., multi-factor authentication in accordance to Grassi et al. [51] raises the security level of
1202 an organization and thus a more substantial contribution to security as its peer authentication approaches. Similar
1203 suggestions about the superior usability of (some) biometric factors like fingerprints compared to passwords are present
1204 in the literature. It is thus possible to aggregate these metrics into an abstract *authentication strength* metric, using its
1205 formula to weigh the different authentication factors. Finally, the identity lifecycle and quality have a close connection.
1206

1207
1208 The policy lifecycle includes metrics assessing the general adaptability of policies, failures to translate policies to
1209 actual ACPs as bypasses, and constraints for policy adaption. The policy lifecycle closely connects with policy quality.
1210 Due to their formulation, bypass, adaptability, and constraints compete with each other, meaning that a maximized
1211 adaptability might also indicate issues at bypass or constraints and vice versa.
1212

1213 From a governance and management perspective, the organization's goals and their comprehensive achievement are
1214 in focus. Establishing means and metrics to control this perspective raises short-term costs. However, this control causes
1215 long-term cost-cutting because of the security benefits, achieving smooth operations, staying compliant, and realizing
1216 high quality. From a long-term IT governance perspective, organizations desire a controlled and comprehensive steering,
1217 reflected by the completeness and control metrics. From a short-term IT management perspective, metrics indicate the
1218 actual capabilities to manage an IAM system, reflected by the traceability, size, and hygiene metrics.
1219

1220 General-purpose process quality metrics show feasibility for IAM, as illustrated for a joiner identity lifecycle process.
1221 While specifics apply, IAM process quality is an instance of general process quality. Despite the shown feasibility of
1222 the QEF process quality metrics [60], this relationship implies that other process quality metrics (like [113, 150]) can
1223 reflect IAM process quality. This modularity allows for the replaceability of QEF for IAM process quality metrics to
1224 other frameworks. Furthermore, this perspective has a powerful influence on the operations IAM goal.
1225

1226 Identity quality relates to the identity lifecycle perspective. This quality goal spans the user entity or application-
1227 specific accounts [85] and derives from general data quality as both represent a set of attributes [116]. Current data
1228 quality literature defines a broad range of metrics [24, 117, 156]. While these generally apply for IAM, this study reduced
1229 them to more relevant data quality metrics for identities. E.g., in an IAM context, identity privacy is crucial and thus
1230 considered with a standalone metric. Proper handling of identity quality influences compliance and operation goals.
1231

1232 Policy quality relates to the policy lifecycle and is an independent quality domain, drawing insights from data [83] and
1233 software quality [161, 162]. Moreover, its performance is crucial: permissive authorizations harm compliance and security,
1234 and missing ones hinder operations. Therefore, policy quality is a research interest of dedicated studies [46, 83, 84].
1235 Besides the authorization accuracy and evaluation runtime as vital metrics for policy quality, their maintenance is
1236 challenging and costly because of the ever-changing environment and organizations themselves. Keeping an eye on
1237 maintainability and its related metrics is thus a crucial factor for quick and informed adjustments whenever required,
1238 effectively benefiting all IAM goals. In that sense, simplicity is the key to maintenance and policy quality.
1239

1240 Software quality is the basis for IAM technology. While hardware quality also influences IAM technology quality, it
1241 is only relevant for specific cases, like dedicated (biometric) authentication devices. This study thus mapped general-
1242 purpose software quality metrics of the ISO25010 standard [76] to IAM. This successful mapping has three implications:
1243 First, IAM depends on general software quality. Second, the mapping highlights special considerations for IAM. Third,
1244 the IAM software quality metrics based on ISO205010 are modular and replaceable by other software quality metric
1245

1246
1247
1248 Manuscript submitted to ACM

1249 frameworks, like [19, 34, 50, 98]. While short-term cost-cutting is negatively affected for achieving high-quality IAM
1250 software, it pays off long-term as other IAM goals are positively affected.

1252 6.2 Relevance of IAM goals

1253 IAM goals like seen in Figure 3 realize IAM strategy [42] in an economic context. IAM is thus a tool for cost-reduction
1254 and quality assurance [110] with further sub-goals [42, 67]. For effective strategic alignment of IAM, IAM metrics require
1255 consideration of IAM goals. This has implications discussed in this section, including their interwoven requirements,
1256 the danger of snake oil, a focus on the big picture, and equilibria within IAM metrics.

1257 IAM goals originate from *interwoven requirements*. These goals thus compete with each other. Access Reviews are
1258 an example of this [52, 78]: While necessary for compliance, inquired domain experts face a time-consuming task
1259 (harming the cost reduction goal), interrupting them from their usual responsibilities (hurting the operation goal). Due
1260 to this pressure, a domain expert might desire an immediate treatment of the Access Review, leading to relatively quick
1261 but unqualified decision-making. Thus, at the expense of security, a domain expert might *just accept* the presented
1262 authorizations since a revoke might cause harm to operations due to missing authorizations. At least this Access Review
1263 fulfills the compliance requirement since the domain expert executed the periodic check. Enforcing the interwoven IAM
1264 goals too strictly thus might set false incentives. Therefore, balanced approaches require consideration for other IAM
1265 goals and aspects. For example, continuously evaluating a broad collection of metrics like those proposed by this work
1266 is helpful. Compliance is another factor for the interwovenness of IAM goals. Compliance can originate from internal
1267 requirements, but external regulative authorities often impose compliance. This dependency has three implications:
1268 First, multiple compliance frameworks might apply simultaneously. For example, an organization might need to comply
1269 with internal goals while also fulfilling the ISO27001 [77] and SOX [149] since it is available on the US public board.
1270 Another organization might operate in the healthcare sector, which requires additional compliance with HIPAA [148].
1271 Despite the similarity of the compliance frameworks, a careful evaluation of the practical requirements is necessary.
1272 Second, the external dependency on compliance is vulnerable to external changes. A new compliance regulation thus
1273 might come into effect, or a present one might receive an update. For example, the GDPR [37] caused some insecurities
1274 in its requirements after coming into effect [28]. Another example of an update is the ISO27001:2013 which received an
1275 update in 2022 (ISO27001:2022). Third, unlike other goals, compliance does not necessarily require a timeless or valid
1276 argumentation but somewhat situational or political reasoning. As in the Access Review example, a compliance-relevant
1277 procedure can make sense in some environments but only sometimes in all. Another example is SOX [149], which came
1278 into effect after major accounting scandals (Enron, WorldCom, etc.), highlighting situational reasoning.

1279 IAM systems are complex constructs facing an ever-changing environment, which is even hard for experts to
1280 understand. The presence of *snake oil* IAM metrics is thus no surprise as (promoted) measurements do not always make
1281 sense regarding IAM goals. For example, a standalone measurement *number of roles* is easily obtainable, visualizable,
1282 and pretends an overview but has a limited expressiveness for IAM. Larger organizations often use heterogenous policy
1283 types [82], different role types (business vs. IT roles), etc. Additionally, the sheer number of roles does not consider
1284 assignments like the more expressive role coverage [84]. Furthermore, the grounding of the role number regarding
1285 IAM goals remains unclear. For this reason, the IAM metric collection of this work bases its metrics on (strategic)
1286 IAM goals. The metrics decouple themselves from measurements to depict more abstract concepts measured by the
1287 measurements. This decoupling and relation to (strategic) IAM goals have a few implications: First, IAM metrics require
1288 a relationship to IAM goals derived from organizational strategy. An IAM metric thus needs to show its impact and
1289 alignment to organizational strategy. Otherwise, the IAM metric's justification is dubious. Second, an IAM metric's

1290
1291
1292
1293
1294
1295
1296
1297
1298
1299
1300

Manuscript submitted to ACM

1301 concept requires decoupling from the actual measurements or underlying data. For example, IAM systems include
1302 organization-specific data and interpretation. The IAM metric concept thus does not change because of an organization
1303 or system. Organizational fit is rather achieved by its measurements.
1304

1305 A grounding of IAM metrics on IAM goals allows for keeping track of the *big picture* since the IAM metrics become
1306 comparable regarding their impact on IAM goals. At first glance, it seems reasonable to investigate the effects of the
1307 IAM metrics in direct comparison. However, regarding the organizational strategic alignment, comparing their impact
1308 on IAM goals is more worthwhile. Comparing IAM goals gives a clearer perspective on the organizational big picture
1309 and strategic alignment. For example, an overdone password strength might raise security at the expense of operations
1310 as performance breaks down. While their effect on each other might be relevant for analytical purposes, balanced
1311 security and operations are the crucial focus for strategic alignment. Actions for restoring the IAM goal balance thus
1312 should focus on achieving the IAM goals again, not on particular IAM metrics, despite their hint at potential issues.
1313

1314 Finally, this study indicates *equilibria within IAM metrics* regarding relationships to IAM goals, highlighting their
1315 real-world issues and research opportunities. An illustrative example is the password strength metric, standing between
1316 security and operations. Morris and Thompson [101] already discussed this equilibrium back in the seventies, despite
1317 that it is still an ongoing research topic [20, 53, 79]. An unbalanced increase of password security parameters contributing
1318 to security harms operations and vice versa. For example, longer, more complex, and more secure passwords for each
1319 service are hard to remember and thus hamper operations but increase security. Conversely, more simplistic passwords
1320 are vulnerable to cracking attempts and thus hamper security but increase operations. These equilibria are areas of
1321 tension since organizations need to find a balance between the two IAM goals affected by this metric. Furthermore, this
1322 study highlights these equilibria whenever an IAM metric in a Table of Section 5 has a positive and negative relationship
1323 towards an IAM goal. These equilibria are an issue and an opportunity for advances in academia and practice.
1324
1325
1326
1327

1328 6.3 Understandability

1329 An essential aspect of this study is its understandability. Various requirements influence IAM reflected by the 7
1330 perspectives applied in this study. A limitation to 7 perspectives is intended as it fosters comprehension. Psychological
1331 research suggests that average humans simultaneously comprehend max. 7 ± 2 [96] or even better only 4 items [29]. This
1332 study incorporates this limitation by restricting itself to 7 perspectives, which contain at maximum 7 ± 2 but aim for 4
1333 IAM metrics to remain comprehensive. Picking the right abstraction level is thus a key to achieving understandability.
1334

1335 More than choosing an understandable abstraction level is required since large amounts of potential metrics and
1336 measurements are present. Metrics need aggregation and provide drill-down functionalities [18]. Aggregation allows a
1337 reduction of the presented metrics. This reduction is necessary to achieve an understandable amount of metrics (e.g.,
1338 7 ± 2). A drawback of this approach is that specific metrics merge. These metrics are no longer available as standalone
1339 metrics but as components of more abstract metrics. Furthermore, aggregation implies an extreme situation for which a
1340 few metrics describe an organization's IAM comparable to other organizations. This extreme is a maturity model derived
1341 from the collected metrics, allowing for comparison of real-world IAM approaches [138]. Opposite to aggregation is
1342 the drill-down functionality. The drill-down allows decomposing a metric into its sub-metrics or measurements. This
1343 decomposing enables analysts to identify and tackle the root causes of failing target values for IAM goals.
1344
1345

1346 Finally, the audience for IAM metrics needs consideration. Since organizations require flexibility, the presented
1347 IAM metrics are no fixed set but a suggestion of valid concepts. This suggestion allows for configurations on the
1348 utilized measurements, formulas (e.g., for ranking and weighting [136]), or even an exclusion of specific metrics or
1349 customized metrics for the organization. This customization ensures the necessary fine-tuning for specific organizational
1350
1351
1352

Manuscript submitted to ACM

1353 environments, fostering a tailored, comprehensive, and aligned IAM metric collection. Furthermore, various audiences
1354 require specific considerations. For example, policy administrators are likely more interested in policy lifecycle and
1355 quality metrics. Compared with that, a staff with a non-technical skill set might be less interested in software quality
1356 metrics. Careful considerations for the specific IAM audiences are thus required to leverage further comprehension.
1357

1358 6.4 Limitations

1359 One limiting aspect of this study is economic feasibility. While database queries might cost-efficiently realize some
1360 measurements (assuming data availability), other measurements might include costly surveys or interviews. Databases
1361 containing the required data need implementation and maintenance. While this economic feasibility is discussed for
1362 the data quality realm in greater detail as for IT security, it is a relevant aspect for both [61]. In summary, the costs for
1363 IAM metrics thus cannot exceed the expected savings by using them leading to more informed decisions down the line.
1364 For example, an organization-wide daily user satisfaction study thus might give insights but also violate this economic
1365 feasibility. However, not running the study may lead to wrong decisions, also causing costs.
1366

1367 Another limiting factor of this study is its necessary and inherent subjectivity. We want to emphasize that the
1368 proposed collection of IAM metrics is not *the single optimal* solution but rather *one valid* solution among many valid
1369 solutions. This solution space implies that other solutions for a comprehensive IAM collection are also feasible and
1370 valid. Like this study, the chosen combination of IAM metrics might base on objective criteria: a solid literature review,
1371 a grounding on IAM goals, limiting their amount for the sake of understandability, or utilizing present collections
1372 of related metrics (like done for software quality). However, this and other identified IAM metric collections would
1373 still be subjective to some degree. This subjectivity is based on the composition of metrics and their flexibility. For
1374 some organizations, measurements or even metrics might not be applicable. For example, a biometric strength metric
1375 is unnecessary if an organization does not use biometric authentication. Additionally, organizations might need to
1376 cover specific needs, forcing them only to use specific, adapted, or newly created measurements or metrics to cover
1377 their context-specific and personal needs [18]. In this sense, this collection of IAM metrics can serve as a suggestion on
1378 appropriate and grounded IAM metrics. However, they need a fitting for the applied context.
1379

1380 7 RELATED WORK

1381 The presented collection of IAM metrics complements related work covering aligning and integrating metrics in an
1382 organizational context, including their goals and strategy. The consequences for misaligned metrics, goals, and strategies
1383 are various [11, 62]. E.g, unknown organizational strategy on the project level, metrics with a wrong or even without
1384 impact on organizational goals, etc. Therefore, it is evident to avoid misalignments for effective decision-making.
1385

1386 Several approaches for integrating metrics in an organizational context are published. E.g., the renowned Goal Question
1387 Metric (GQM) approach [8] uses a cascade to derive expressive metrics based on goals. It starts with conceptual
1388 goals, later formulated in greater detail with questions. These questions then allow for a quantitative assessment of the
1389 goal by using metrics. This approach can also be extended with the GQM+Strategies approach [9], which underlines the
1390 need to align metrics with goals and derive goals from organizational strategies. Similar approaches are the Balanced
1391 Scorecard (BSC) [80] and the Model, Measure, Manage Paradigm (M3P) [108]. The BSC compactly presents complex
1392 information at a glance for management audiences. The perspectives of customer, financial, internal business, and
1393 innovation and learning connect goals to measures. This approach allows for a strategic perspective on metrics. The
1394 M3P framework also puts the metrics in a more holistic context of an organization. It highlights identifying and
1395 understanding strategies and organizational goals to yield effective metrics and measurements.
1396

1400
1401
1402
1403
1404 Manuscript submitted to ACM

1405 From an IT governance perspective, Control Objectives for Information and Related Technologies (COBIT) [72]
1406 and Information Technology Infrastructure Library (ITIL) [146] are well-known complementary frameworks. These
1407 frameworks comprehensively align organizational strategies and goals with controls and measurements. While widely
1408 adopted in the industry and still improved with their versions COBIT 2019 and ITIL 4, these frameworks lack the
1409 simplicity of the GQM or the BSC. However, both COBIT and ITIL cascade strategies and goals towards metrics.

1411 These frameworks for integrating metrics thus benefit from (IAM) metric collections, like presented by this work.
1412 Similar yet less comprehensive IAM metric collections are discussed in the literature. Hummer et al. [67] derive
1413 performance indicators from IAM goals using the GQM method and evaluate the results by an expert survey. A
1414 limitation of this contribution is its lacking comprehensiveness, suggested as future work. Elimity [35] derives eight
1415 IAM metrics with measurements from ISO27001 and connects them to IAM goals, alike the ones argued by Hummer et
1416 al. [67]. Additionally, the authors suggest a feedback loop based on the metrics. A limitation of this contribution is a
1417 missing peer review, and its sole focus is on the ISO27001. Hu et al. [65, 66] contribute via a NIST publications with
1418 access control metrics. An XACML based architecture, ontology, and responsible principals classify measurements
1419 to metrics. While this contribution is more comprehensive than the other contributions covered by this related work
1420 section, it lacks the coverage of quality factors. The authors also consider composing the metrics for custom use cases.

1424 With this related work, system architects can implement IAM metrics in organizations. As suggested by Kern et al.
1425 [82], governance stakeholders define and prioritize IAM goals aligned with organizational strategy. With this scope,
1426 implementors set their objectives and allocate suitable IAM metrics for relevant audiences. Additionally, implementors
1427 must fit the IAM metrics to their organization by adjusting the weighting and utilization of context-specific measure-
1428 ments. IAM metrics then monitor the effects of used tools and taken actions. Stakeholders, like the implementors or
1429 governance, receive IAM metric reports regularly, e.g., on a dedicated IAM governance dashboard [120].

1431

1432 8 CONCLUSION

1433

1434 Our work identifies 43 IAM metrics spanning over seven perspectives while limiting itself to understandable dimensions,
1435 metrics with an impact on IAM goals, and relevance in literature. These IAM metrics enable organizations to keep a
1436 holistic overview of essential IAM topics and trace an effective utilization of their security budget for IAM. To the best
1437 of our knowledge, this work is the most comprehensive scientific study published for IAM metrics and their alignment
1438 to strategic IAM goals. These IAM metrics based on IAM goals provide an accessible and broad overview for IAM
1439 while staying understandable. This work also includes guidance for implementing these IAM metrics in organizational
1440 strategy and hints at practical features like drill-down and aggregation.

1442

1443 We identify multiple directions for future work. First, any time an IAM metric has a negative and positive effect
1444 on at least two IAM goals, a research opportunity is present. Future work targeting these IAM metrics thus needs
1445 careful consideration and balancing for these competing IAM goals. Another direction for future work is applying this
1446 collection of IAM metrics for various organizations. Such a study can compare different organizations and might be
1447 the basis for determining thresholds and labels of the IAM metrics indicating an IAM maturity model. At the time of
1448 submission, we are working on a prototype that manages (custom) IAM metrics, their alignment with IAM goals and
1449 strategy, and their visualization for (custom) IAM audiences.

1451

1452 ACKNOWLEDGMENTS

1453

1454 This work was supported by the German Federal Ministry of Education and Research with the project DEVISE. We also
1455 want to thank our students: Marco Rauchecker, Johanna Kopp, Nikola Strobel, and Mathis Müller.

1456

Manuscript submitted to ACM

REFERENCES

- [1] Aymen Akremi and Mohsen Rouached. 2021. A comprehensive and holistic knowledge model for cloud privacy protection. *The Journal of Supercomputing* 77, 8 (01 Aug 2021), 7956–7988. <https://doi.org/10.1007/s11227-020-03594-3>
- [2] Andrew van der Stock, Brian Glas, Neil Smithline, Torsten Giegler. 2021. *OWASP Top 10:2021*. Technical Report. Open Web Application Security Project (OWASP). <https://owasp.org/www-project-top-ten/>
- [3] Patricia Arias-Cabarcos, Florina Almenárez-Mendoza, Andrés Marín-López, Daniel Díaz-Sánchez, and Rosa Sánchez-Guerrero. 2012. A Metric-Based Approach to Assess Risk for “On Cloud” Federated Identity Management. *Journal of Network and Systems Management* 20, 4 (01 Dec 2012), 513–533. <https://doi.org/10.1007/s10922-012-9244-2>
- [4] Patricia Arias-Cabarcos, Christian Krupitzer, and Christian Becker. 2019. A Survey on Adaptive Authentication. *ACM Comput. Surv.* 52, 4, Article 80 (sep 2019), 30 pages. <https://doi.org/10.1145/3336117>
- [5] Rashid Arsalan, Lateef Mehreen, Kaur Balbir, O.P. Aggarwal, Hamid Sajad, and Gupta Neeraj. 2013. Biometric Finger Print Identification Is It a Reliable Tool or Not? *Journal of Indian Academy of Forensic Medicine* 35, 2 (2013), 109–112.
- [6] Aaron Bangor, Philip T. Kortum, and James T. Miller. 2008. An Empirical Evaluation of the System Usability Scale. *International Journal of Human-Computer Interaction* 24, 6 (2008), 574–594. <https://doi.org/10.1080/10447310802205776> arXiv:<https://doi.org/10.1080/10447310802205776>
- [7] Basel Committee on Banking Supervision. 2011. Basel III: A global regulatory framework for more resilient banks and banking systems.
- [8] Victor Basili, Gianluigi Caliera, and H. Dieter Rombach. 1994. The Goal Question Metric Approach. *Encyclopedia of Software Engineering* 1, 1 (1994), 528–532.
- [9] Victor Basili, Jens Heidrich, Mikael Lindvall, Jurgen Munch, Myrna Regardie, and Adam Trendowicz. 2007. GQM+ Strategies – Aligning Business Strategies with Software Measurement. In *First International Symposium on Empirical Software Engineering and Measurement (ESEM 2007)*. IEEE, Madrid, Spain, 488–490. <https://doi.org/10.1109/ESEM.2007.66>
- [10] Thomas Baumer, Mathis Müller, and Günther Pernul. 2023. System for Cross-Domain Identity Management (SCIM): Survey and Enhancement With RBAC. *IEEE Access* 11 (2023), 86872–86894. <https://doi.org/10.1109/ACCESS.2023.3304270>
- [11] Shirley A. Becker and Mitchell L. Bostelman. 1999. Aligning strategic and project measurement systems. *IEEE Software* 16, 3 (May 1999), 46–51. <https://doi.org/10.1109/52.765786>
- [12] Matthias Beckerle and Leonardo A. Martucci. 2013. Formal Definitions for Usable Access Control Rule Sets from Goals to Metrics. In *Proceedings of the Ninth Symposium on Usable Privacy and Security (Newcastle, United Kingdom) (SOUPS '13)*. Association for Computing Machinery, New York, NY, USA, Article 2, 11 pages. <https://doi.org/10.1145/2501604.2501606>
- [13] Marios Belk, Christos Fidas, Panagiotis Germanakos, and George Samaras. 2015. Do human cognitive differences in information processing affect preference and performance of CAPTCHA? *International Journal of Human-Computer Studies* 84 (2015), 1–18. <https://doi.org/10.1016/j.ijhcs.2015.07.002>
- [14] Marios Belk, Panagiotis Germanakos, Christos Fidas, and George Samaras. 2013. Studying the Effect of Human Cognition on User Authentication Tasks. In *User Modeling, Adaptation, and Personalization*, Sandra Carberry, Stephan Weibelzahl, Alessandro Micarelli, and Giovanni Semeraro (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 102–113. https://doi.org/10.1007/978-3-642-38844-6_9
- [15] Yolanta Beres, Marco Casassa Mont, Jonathan Griffin, and Simon Shiu. 2009. Using Security Metrics Coupled with Predictive Modeling and Simulation to Assess Security Processes. In *Proceedings of the 2009 3rd International Symposium on Empirical Software Engineering and Measurement (ESEM '09)*. IEEE Computer Society, USA, 564–573. <https://doi.org/10.1109/ESEM.2009.5314213>
- [16] Beyond Identity. 2022. Former employees admit to using continued account access to harm previous employers. <https://www.beyondidentity.com/blog/great-resignation-impact-on-company-security>
- [17] Debnath Bhattacharyya, Rahul Ranjan, Farkhod Alisherov, Minkyu Choi, et al. 2009. Biometric authentication: A review. *International Journal of u-and e-Service, Science and Technology* 2, 3 (2009), 13–28.
- [18] Paul Black, Karen Scarfone, and Murugiah Souppaya. 2009. *Cyber Security Metrics and Measures*. John Wiley & Sons, Inc., Hoboken, NJ, Nebraska. https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=51292
- [19] B. W. Boehm, J. R. Brown, and M. Lipow. 1976. Quantitative Evaluation of Software Quality. In *Proceedings of the 2nd International Conference on Software Engineering (San Francisco, California, USA) (ICSE '76)*. IEEE Computer Society Press, Washington, DC, USA, 592–605. <https://doi.org/10.5555/800253.807736>
- [20] Christina Braz, Ahmed Seffah, and David M'Raihi. 2007. Designing a Trade-Off Between Usability and Security: A Metrics Based-Model. In *Human-Computer Interaction – INTERACT 2007*, Cécilia Baranauskas, Philippe Palanque, Julio Abascal, and Simone Diniz Junqueira Barbosa (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 114–126. https://doi.org/10.1007/978-3-540-74800-7_9
- [21] John Brooke. 1996. SUS: A Quick and Dirty Usability Scale. In *Usability Evaluation In Industry*. Taylor and Francis, London, United Kingdom, 207–212. <https://doi.org/10.1201/9781498710411-35>
- [22] John Brooke. 2013. SUS: A Retrospective. *J. Usability Studies* 8, 2 (feb 2013), 29–40.
- [23] Elizabeth Chew, Marianne Swanson, Kevin Stine, Nadya Bartol, Anthony Brown, and Will Robinson. 2008. *Performance measurement guide for information security*. Technical Report. National Institute of Standards and Technology, Gaithersburg, MD.
- [24] Corinna Cichy and Stefan Rass. 2019. An Overview of Data Quality Frameworks. *IEEE Access* 7 (2019), 24634–24648. <https://doi.org/10.1109/ACCESS.2019.2899751>

30

Baumer et al.

- 1509 [25] Sebastian Clauß. 2006. A Framework for Quantification of Linkability Within a Privacy-Enhancing Identity Management System. In *Emerging*
1510 *Trends in Information and Communication Security*, Günter Müller (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 191–205. [https://doi.org/](https://doi.org/10.1007/11766155_14)
1511 [10.1007/11766155_14](https://doi.org/10.1007/11766155_14)
- 1512 [26] Sebastian Clauß and Stefan Schiffner. 2006. Structuring Anonymity Metrics. In *Proceedings of the Second ACM Workshop on Digital Identity*
1513 *Management (Alexandria, Virginia, USA) (DIM '06)*. Association for Computing Machinery, New York, NY, USA, 55–62. [https://doi.org/10.1145/](https://doi.org/10.1145/1179529.1179539)
1514 [1179529.1179539](https://doi.org/10.1145/1179529.1179539)
- 1515 [27] Gary Cole. 2005. *OASIS Service Provisioning Markup Language (SPML) Version 2*. Committee Draft 1.0 pstc-spml2-cd-01. OASIS. [https://docs.oasis-](https://docs.oasis-open.org/provision/spml-2.0-cd-01/pstc-spml2-cd-01.pdf)
1516 [open.org/provision/spml-2.0-cd-01/pstc-spml2-cd-01.pdf](https://docs.oasis-open.org/provision/spml-2.0-cd-01/pstc-spml2-cd-01.pdf) Committee Draft 1.0.
- 1517 [28] Alison Cool. 2018. Europe’s data protection law is a big, confusing mess. [https://www.nytimes.com/2018/05/15/opinion/gdpr-europe-data-](https://www.nytimes.com/2018/05/15/opinion/gdpr-europe-data-protection.html)
1518 [protection.html](https://www.nytimes.com/2018/05/15/opinion/gdpr-europe-data-protection.html)
- 1519 [29] Nelson Cowan. 2001. The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences* 24,
1 (2001), 87–114. <https://doi.org/10.1017/S0140525X01003922>
- 1520 [30] Jon Currey, Robbie McKinstry, Armon Dadgar, and Mark Gritter. 2020. Informed Privilege-Complexity Trade-Offs in RBAC Configuration. In
1521 *Proceedings of the 25th ACM Symposium on Access Control Models and Technologies (Barcelona, Spain) (SACMAT '20)*. Association for Computing
1522 Machinery, New York, NY, USA, 119–130. <https://doi.org/10.1145/3381991.3395597>
- 1523 [31] Xavier de Carné de Carnavalet and Mohammad Mannan. 2014. From Very Weak to Very Strong: Analyzing Password-Strength Meters. In *Network*
1524 *and Distributed System Security Symposium NDSS 2014*. Internet Society, San Diego, California, 1–16. <https://doi.org/10.14722/ndss.2014.23268>
- 1525 [32] V. Neelaya Dhatchayani and V.S. Shankar Sriram. 2014. Trust aware identity management for cloud computing. *International Journal of Information*
1526 *and Communication Technology* 6, 3/4 (2014), 369. <https://doi.org/10.1504/ijict.2014.063220>
- 1527 [33] Rainer Diesch and Helmut Krömer. 2020. SoK: Linking Information Security Metrics to Management Success Factors. In *Proceedings of the 15th*
1528 *International Conference on Availability, Reliability and Security (Virtual Event, Ireland) (ARES '20)*. Association for Computing Machinery, New
1529 York, NY, USA, Article 98, 10 pages. <https://doi.org/10.1145/3407023.3407059>
- 1530 [34] R. Geoff Dromey. 1996. Cornering the Chimera [software quality]. *IEEE Software* 13, 1 (Jan 1996), 33–43. <https://doi.org/10.1109/52.476284>
- 1531 [35] Elimity. 2020. KPI-driven approach to Identity & Access Management - A guide to maximize the value of IAM. [https://go.elimity.com/downloads/KPI-](https://go.elimity.com/downloads/KPI-driven-approach-to-IAM-guide.pdf)
1532 [driven-approach-to-IAM-guide.pdf](https://go.elimity.com/downloads/KPI-driven-approach-to-IAM-guide.pdf)
- 1533 [36] Aaron Elliott and Scott Knight. 2010. Role explosion: Acknowledging the problem.. In *Software Engineering research and practice*. CSREA Press, Las
1534 Vegas, Nevada, 349–355.
- 1535 [37] European Commission. 2016. General Data Protection Regulation.
- 1536 [38] David F. Ferraiolo, Ravi Sandhu, Serban Gavrilă, D. Richard Kuhn, and Ramaswamy Chandramouli. 2001. Proposed NIST Standard for Role-Based
1537 Access Control. *ACM Trans. Inf. Syst. Secur.* 4, 3 (aug 2001), 224–274. <https://doi.org/10.1145/501978.501980>
- 1538 [39] FIRST. 2019. CVSS v3.1 Specification Document. <https://www.first.org/cvss/specification-document>
- 1539 [40] Simone Fischer-Hübner, Cristina Alcaraz, Afonso Ferreira, Carmen Fernandez-Gago, Javier Lopez, Evangelos Markatos, Lejla Islami, and Mahdi
1540 Akil. 2021. Stakeholder perspectives and requirements on cybersecurity in Europe. *Journal of Information Security and Applications* 61 (2021),
1541 102916. <https://doi.org/10.1016/j.jisa.2021.102916>
- 1542 [41] Ludwig Fuchs, Michael Kunz, and Günther Pernul. 2014. Role Model Optimization For Secure Role-based Identity Management. In *European*
1543 *Conference on Information Systems (ECIS)*. AIS, Tel Aviv, Israel, 1–15. <https://epub.uni-regensburg.de/30394/>
- 1544 [42] Ludwig Fuchs and Gunther Pernul. 2007. Supporting Compliant and Secure User Handling - A Structured Approach for In-House Identity
1545 Management. In *The Second International Conference on Availability, Reliability and Security (ARES'07)*. IEEE, Vienna, Austria, 374–384. <https://doi.org/10.1109/ARES.2007.145>
- 1546 [43] Ludwig Fuchs and Günther Pernul. 2008. HyDro – Hybrid Development of Roles. In *Information Systems Security*, R. Sekar and Arun K. Pujari
1547 (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 287–302.
- 1548 [44] Ludwig Fuchs and Günther Pernul. 2008. proROLE: A Process-oriented Lifecycle Model for Role Systems. In *Proceedings of the 16th European*
1549 *Conference on Information Systems (ECIS 2008), Galway, Ireland, June 9-11, 2008*. Springer, Berlin, 1322–1333. <https://aisel.aisnet.org/ecis2008/111>
- 1550 [45] Ludwig Fuchs and Günther Pernul. 2010. Reducing the Risk of Insider Misuse by Revising Identity Management and UserAccount Data. *Journal of*
1551 *Wireless Mobile Networks, Ubiquitous Computing and Dependable Applications (JoWUA)* 1, 1 (2010), 14–28. [https://doi.org/10.22667/JoWUA.2010.](https://doi.org/10.22667/JoWUA.2010.06.31.014)
1552 [06.31.014](https://doi.org/10.22667/JoWUA.2010.06.31.014)
- 1553 [46] Ludwig Fuchs, Günther Pernul, and Ravi Sandhu. 2011. Roles in information security – A survey and classification of the research area. *Computers*
1554 *& Security* 30, 8 (2011), 748–769. <https://doi.org/10.1016/j.cose.2011.08.002>
- 1555 [47] Inc. Gartner. 17.05.2021. Gartner Forecasts Worldwide Security and Risk Management Spending to Exceed \$150 Billion in 2021. <https://www.gartner.com/en/newsroom/press-releases/2021-05-17-gartner-forecasts-worldwide-security-and-risk-management>
- 1556 [48] Martin D. Gibbs. 2010. Biometrics: Body Odor Authentication Perception and Acceptance. *SIGCAS Comput. Soc.* 40, 4 (dec 2010), 16–24.
1557 <https://doi.org/10.1145/1929609.1929612>
- 1558 [49] Hidehito Gomi. 2011. An Authentication Trust Metric for Federated Identity Management Systems. In *Security and Trust Management*, Jorge Cuellar,
1559 Javier Lopez, Gilles Barthe, and Alexander Pretschner (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 116–131. [https://doi.org/10.1007/978-](https://doi.org/10.1007/978-3-642-22444-7_8)
1560 [3-642-22444-7_8](https://doi.org/10.1007/978-3-642-22444-7_8)

Manuscript submitted to ACM

Identity and Access Management Metrics

31

- 1561 [50] Robert B. Grady. 1992. *Practical Software Metrics for Project Management and Process Improvement*. Prentice-Hall, Inc., USA. <https://doi.org/10.5555/140207>
- 1562
- 1563 [51] Paul A. Grassi, James L. Fenton, Elaine M. Newton, Ray A. Perlner, Andrew R. Regenscheid, William E. Burr, Justin P. Richer, Naomi B. Lefkovitz, Jamie M. Danker, Yee-Yin Choong, Kristen K. Greene, and Mary F. Theofanos. 2017. *Digital identity guidelines: authentication and lifecycle management*. Technical Report. National Institute of Standards and Technology. <https://doi.org/10.6028/nist.sp.800-63b>
- 1564
- 1565 [52] Sebastian Groll, Sascha Kern, Ludwig Fuchs, and Günther Pernul. 2021. Monitoring Access Reviews by Crowd Labelling. In *Trust, Privacy and Security in Digital Business*, Simone Fischer-Hübner, Costas Lambrinoudakis, Gabriele Kotsis, A. Min Tjoa, and Ismail Khalil (Eds.). Springer International Publishing, Cham, 3–17.
- 1566
- 1567
- 1568 [53] Hana Habib, Pardis Emami Naeini, Summer Devlin, Maggie Oates, Chelse Swoopes, Lujo Bauer, Nicolas Christin, and Lorrie Faith Cranor. 2018. User Behaviors and Attitudes Under Password Expiration Policies. In *Fourteenth Symposium on Usable Privacy and Security (SOUPS 2018)*. USENIX Association, Baltimore, MD, 13–30. <https://www.usenix.org/conference/soups2018/presentation/habib-password>
- 1569
- 1570
- 1571 [54] Umme Habiba, Abdul Ghafoor Abassi, Rahat Masood, and Muhammad Awais Shibli. 2013. Assessment Criteria for Cloud Identity Management Systems. In *2013 IEEE 19th Pacific Rim International Symposium on Dependable Computing*. IEEE, Vancouver, BC, Canada, 188–195. <https://doi.org/10.1109/PRDC.2013.39>
- 1572
- 1573
- 1574 [55] Umme Habiba, Rahat Masood, Muhammad Awais Shibli, and Muaz A. Niazi. 2014. Cloud identity management security issues & solutions: a taxonomy. *Complex Adaptive Systems Modeling* 2, 1 (11 Nov 2014), 5. <https://doi.org/10.1186/s40294-014-0005-9>
- 1575
- 1576 [56] Kemal Hajdarevic and Pat Allen. 2013. A new method for the identification of proactive information security management system metrics. In *2013 36th International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO)*. IEEE, Opatija, Croatia, 1121–1126.
- 1577
- 1578 [57] John Halamka, Ari Juels, Adam Stubblefield, and Jonathan Westhues. 2006. The Security Implications of VeriChip Cloning. *Journal of the American Medical Informatics Association* 13, 6 (11 2006), 601–607. <https://doi.org/10.1197/jamia.M2143>
- 1579
- 1580 [58] Marit Hansen, Peter Berlich, Jan Camenisch, Sebastian Clauß, Andreas Pfitzmann, and Michael Waidner. 2004. Privacy-enhancing identity management. *Information Security Technical Report* 9, 1 (2004), 35–44. [https://doi.org/10.1016/S1363-4127\(04\)00014-7](https://doi.org/10.1016/S1363-4127(04)00014-7)
- 1581
- 1582 [59] Dick Hardt. 2012. The OAuth 2.0 Authorization Framework. RFC 6749. <https://doi.org/10.17487/RFC6749>
- 1583
- 1584 [60] Farideh Heidari and Pericles Loucoupoulos. 2014. Quality evaluation framework (QEF): Modeling and evaluating quality of business processes. *International Journal of Accounting Information Systems* 15, 3 (2014), 193–223. <https://doi.org/10.1016/j.iaacinf.2013.09.002>
- 1585
- 1586 [61] Bernd Heinrich, Diana Hristova, Mathias Klier, Alexander Schiller, and Michael Szubartowicz. 2018. Requirements for Data Quality Metrics. *J. Data and Information Quality* 9, 2, Article 12 (jan 2018), 32 pages. <https://doi.org/10.1145/3148238>
- 1587
- 1588 [62] J. C. Henderson and H. Venkatraman. 1993. Strategic alignment: Leveraging information technology for transforming organizations. *IBM Systems Journal* 32, 1 (1993), 472–484. <https://doi.org/10.1147/sj.382.0472>
- 1589
- 1590 [63] Markus Hornsteiner, Sebastian Groll, and Alexander Puchta. 2021. Towards a User-Centric IAM Entitlement Shop - Learnings from the e-Commerce. In *13th International Conference on Security of Information and Networks (Merkez, Turkey) (SIN 2020)*. Association for Computing Machinery, New York, NY, USA, Article 15, 4 pages. <https://doi.org/10.1145/3433174.3433585>
- 1591
- 1592 [64] Vincent C. Hu, David Ferraiolo, Rick Kuhn, Arthur R. Friedman, Alan J. Lang, Margaret M. Cogdell, Adam Schnitzer, Kenneth Sandlin, Robert Miller, Karen Scarfone, et al. 2014. *Guide to attribute based access control (abac) definition and considerations (draft)*. Technical Report. National Institute of Standards and Technology. 1–54 pages. <https://doi.org/10.6028/nist.sp.800-162>
- 1593
- 1594 [65] Vincent C. Hu, David F. Ferraiolo, and D. Rick Kuhn. 2006. *Assessment of access control systems*. Technical Report. National Institute of Standards and Technology. <https://doi.org/10.6028/nist.ir.7316>
- 1595
- 1596 [66] Vincent C. Hu and Karen Scarfone. 2012. *Guidelines for Access Control System Evaluation Metrics*. Technical Report. National Institute of Standards and Technology.
- 1597
- 1598 [67] Matthias Hummer, Sebastian Groll, Michael Kunz, Ludwig Fuchs, and Günther Pernul. 2018. Measuring Identity and Access Management Performance - An Expert Survey on Possible Performance Indicators. In *Proceedings of the 4th International Conference on Information Systems Security and Privacy*. SCITEPRESS - Science and Technology Publications, Funchal, Madeira, Portugal, 233–240. <https://doi.org/10.5220/0006557702330240>
- 1599
- 1600 [68] Matthias Hummer, Michael Kunz, Michael Netter, Ludwig Fuchs, and Günther Pernul. 2016. Adaptive identity and access management—contextual data based policies. *EURASIP Journal on Information Security* 2016, 1 (15 Aug 2016), 19. <https://doi.org/10.1186/s13635-016-0043-2>
- 1601
- 1602 [69] Phil Hunt, Kelly Grizzle, Morteza Ansari, Erik Wahlstroem, and Chuck Mortimore. 2015. System for Cross-domain Identity Management: Protocol. RFC 7644. <https://doi.org/10.17487/RFC7644>
- 1603
- 1604 [70] Phil Hunt, Kelly Grizzle, Erik Wahlstroem, and Chuck Mortimore. 2015. System for Cross-domain Identity Management: Core Schema. RFC 7643. <https://doi.org/10.17487/RFC7643>
- 1605
- 1606 [71] Thorsten Höllrigl, Frank Schell, Sebastian Suelmann, and Hannes Hartenstein. 2008. Towards Systematic Engineering of Service-Oriented Access Control in Federated Environments. In *2008 IEEE Congress on Services Part II (services-2 2008)*. IEEE, Beijing, China, 104–111. <https://doi.org/10.1109/SERVICES-2.2008.24>
- 1607
- 1608 [72] ISACA. 2018. Introducing COBIT 2019 Overview. <https://www.isaca.org/resources/cobit>
- 1609
- 1610 [73] Shareeful Islam and Paolo Falcarin. 2011. Measuring security requirements for software security. In *2011 IEEE 10th International Conference on Cybernetic Intelligent Systems (CIS)*. IEEE, London, UK, 70–75. <https://doi.org/10.1109/CIS.2011.6169137>
- 1611
- 1612 [74] ISO 14223. 2018. *Radiofrequency identification of animals*. Standard. International Organization for Standardization.

Manuscript submitted to ACM

32

Baumer et al.

- 1613 [75] ISO 24760. 2019. *IT Security and Privacy – A framework for identity management*. Standard. International Organization for Standardization.
- 1614 [76] ISO 25010. 2011. *Systems and software engineering – Systems and software Quality Requirements and Evaluation (SQuaRE) – System and software*
- 1615 *quality models*. Standard. International Organization for Standardization.
- 1616 [77] ISO 27001. 2013. *Information technology – Security techniques – Information security management systems – Requirements*. Standard. International
- 1617 Organization for Standardization.
- 1618 [78] Pooya Jaferian, Hootan Rashtian, and Konstantin Beznosov. 2014. To Authorize or Not Authorize: Helping Users Review Access Policies in
- 1619 Organizations. In *Proceedings of the Tenth USENIX Conference on Usable Privacy and Security* (Menlo Park, CA) (SOUPS '14). USENIX Association,
- 1620 USA, 301–320. <https://doi.org/10.5555/3235838.3235865>
- 1621 [79] Gokul Chettoor Jayakrishnan, Gangadhara Reddy Sirigireddy, Sukanya Vaddepalli, Vijayanand Banahatti, Sachin Premsukh Lodha, and
- 1622 Sankalp Suneel Pandit. 2020. Password: A Serious Game to Promote Password Awareness and Diversity in an Enterprise. In *Sixteenth Symposium*
- 1623 *on Usable Privacy and Security (SOUPS 2020)*. USENIX Association, online, 1–18. [https://www.usenix.org/conference/soups2020/presentation/](https://www.usenix.org/conference/soups2020/presentation/jayakrishnan)
- 1624 [jayakrishnan](https://www.usenix.org/conference/soups2020/presentation/jayakrishnan)
- 1625 [80] Robert S. Kaplan and David P. Norton. 1992. The Balanced Scorecard—Measures that Drive Performance. *Harvard Business Review* 1, 1 (1992),
- 1626 71–79. <https://hbr.org/1992/01/the-balanced-scorecard-measures-that-drive-performance-2>
- 1627 [81] Christina Katsini, Marios Belk, Christos Fidas, Nikolaos Avouris, and George Samaras. 2016. Security and Usability in Knowledge-Based User
- 1628 Authentication: A Review. In *Proceedings of the 20th Pan-Hellenic Conference on Informatics* (Patras, Greece) (PCI '16). Association for Computing
- 1629 Machinery, New York, NY, USA, Article 63, 6 pages. <https://doi.org/10.1145/3003733.3003764>
- 1630 [82] Sascha Kern, Thomas Baumer, Ludwig Fuchs, and Günther Pernul. 2023. Maintain High-Quality Access Control Policies: An Academic and Practice-
- 1631 Driven Approach. In *Data and Applications Security and Privacy XXXVII*. Vijayalakshmi Atluri and Anna Lisa Ferrara (Eds.). Springer Nature
- 1632 Switzerland, Cham, 223–242. https://doi.org/10.1007/978-3-031-37586-6_14
- 1633 [83] Sascha Kern, Thomas Baumer, Sebastian Groll, Ludwig Fuchs, and Günther Pernul. 2022. Optimization of Access Control Policies. *Journal of*
- 1634 *Information Security and Applications* 70 (2022), 103301. <https://doi.org/10.1016/j.jjsa.2022.103301>
- 1635 [84] Michael Kunz, Ludwig Fuchs, Michael Netter, and Günther Pernul. 2015. How to Discover High-Quality Roles? A Survey and Dependency Analysis
- 1636 of Quality Criteria in Role Mining. In *Information Systems Security and Privacy*, Olivier Camp, Edgar Weippl, Christophe Bidan, and Esma Aïmeur
- 1637 (Eds.). Springer International Publishing, Cham, 49–67.
- 1638 [85] Michael Kunz, Alexander Puchta, Sebastian Groll, Ludwig Fuchs, and Günther Pernul. 2019. Attribute quality management for dynamic identity
- 1639 and access management. *Journal of Information Security and Applications* 44 (2019), 64–79. <https://doi.org/10.1016/j.jjsa.2018.11.004>
- 1640 [86] Jesse Leskinen. 2012. Evaluation Criteria for Future Identity Management. In *2012 IEEE 11th International Conference on Trust, Security and Privacy*
- 1641 *in Computing and Communications*. IEEE, Liverpool, UK, 801–806. <https://doi.org/10.1109/TrustCom.2012.153>
- 1642 [87] Yair Levy and Timothy J. Ellis. 2006. A Systems Approach to Conduct an Effective Literature Review in Support of Information Systems Research.
- 1643 *Informing Sci. Int. J. an Emerg. Transdiscipl.* 9 (2006), 181–212.
- 1644 [88] James R. Lewis. 2018. The System Usability Scale: Past, Present, and Future. *International Journal of Human-Computer Interaction* 34, 7 (2018),
- 1645 577–590. <https://doi.org/10.1080/10447318.2018.1455307> arXiv:<https://doi.org/10.1080/10447318.2018.1455307>
- 1646 [89] Kepeng Li, Phil Hunt, Bhumip Khasnabish, Anthony Nadalin, and Zachary Zeltsan. 2015. System for Cross-domain Identity Management:
- 1647 Definitions, Overview, Concepts, and Requirements. RFC 7642. <https://doi.org/10.17487/RFC7642>
- 1648 [90] Emil Constantin Lupu. 1998. *A role based framework for distributed systems management*. Ph.D. Dissertation. University of London London,
- 1649 England.
- 1650 [91] Vashek Matyáš and Zdeněk Řiha. 2010. Security of biometric authentication systems. In *2010 International Conference on Computer Information*
- 1651 *Systems and Industrial Management Applications (CISIM)*. IEEE, Krakow, Poland, 19–28. <https://doi.org/10.1109/CISIM.2010.5643698>
- 1652 [92] Peter Mayer, Jan Kirchner, and Melanie Volkamer. 2017. A Second Look at Password Composition Policies in the Wild: Comparing Samples
- 1653 from 2010 and 2016. In *Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017)*. USENIX Association, Santa Clara, CA, 13–28.
- 1654 <https://www.usenix.org/conference/soups2017/technical-sessions/presentation/mayer>
- 1655 [93] Stefan Meier, Ludwig Fuchs, and Günther Pernul. 2013. Managing the Access Grid - A Process View to Minimize Insider Misuse Risks. In *11.*
- 1656 *Internationale Tagung Wirtschaftsinformatik, Leipzig, Germany, February 27 – March 1, 2013*. AIS, Leipzig, Germany, 66. [http://aisel.aisnet.org/](http://aisel.aisnet.org/wi2013/66)
- 1657 [wi2013/66](http://aisel.aisnet.org/wi2013/66)
- 1658 [94] Microsoft Corporation. 2023. Microsoft Digital Defense Report 2022. [https://www.microsoft.com/en-us/security/business/microsoft-digital-](https://www.microsoft.com/en-us/security/business/microsoft-digital-defense-report-2022)
- 1659 [defense-report-2022](https://www.microsoft.com/en-us/security/business/microsoft-digital-defense-report-2022)
- 1660 [95] Martin Mihajlov, Borka Jerman Blazic, and Saso Josimovski. 2011. Quantifying Usability and Security in Authentication. In *2011 IEEE 35th Annual*
- 1661 *Computer Software and Applications Conference*. IEEE, Munich, Germany, 626–629. <https://doi.org/10.1109/COMPSAC.2011.87>
- 1662 [96] George Armitage Miller. 1956. The magical number seven plus or minus two: some limits on our capacity for processing information. *Psychol Rev*
- 1663 63, 2 (March 1956), 81–97.
- 1664 [97] Natalia Miloslavskaya. 2016. Security Operations Centers for Information Security Incident Management. In *2016 IEEE 4th International Conference*
- 1665 *on Future Internet of Things and Cloud (FiCloud)*. IEEE, Vienna, Austria, 131–136. <https://doi.org/10.1109/FiCloud.2016.26>
- 1666 [98] Tsvetelina Mladenova. 2020. Software Quality Metrics – Research, Analysis and Recommendation. In *2020 International Conference Automatics and*
- 1667 *Informatics (ICAI)*. IEEE, Varna, Bulgaria, 1–5. <https://doi.org/10.1109/ICAI50593.2020.9311361>

1668 Manuscript submitted to ACM

- 1665 [99] Ian Molloy, Hong Chen, Tiancheng Li, Qihua Wang, Ninghui Li, Elisa Bertino, Seraphin Calo, and Jorge Lobo. 2008. Mining Roles with Semantic
1666 Meanings. In *Proceedings of the 13th ACM Symposium on Access Control Models and Technologies* (Estes Park, CO, USA) (SACMAT '08). Association
1667 for Computing Machinery, New York, NY, USA, 21–30. <https://doi.org/10.1145/1377836.1377840>
- 1668 [100] Keith L. Monson, Maria Antonia Roberts, Kathryn B. Knorr, Sherine Ali, Stephen B. Meagher, Kevin Biggs, Patti Blume, Donna Brandelli, Albert
1669 Marzioli, Robert Reneau, and Frank Tarasi. 2019. The permanence of friction ridge skin and persistence of friction ridge skin and impressions: A
1670 comprehensive review and new results. *Forensic Science International* 297 (2019), 111–131. <https://doi.org/10.1016/j.forsciint.2019.01.046>
- 1671 [101] Robert Morris and Ken Thompson. 1979. Password Security: A Case History. *Commun. ACM* 22, 11 (nov 1979), 594–597. <https://doi.org/10.1145/359168.359172>
- 1672 [102] Nkosinathi Mpofo and Wynand JC van Staden. 2014. A survey of trust issues constraining the growth of Identity Management-as-a-Service(IdMaaS).
1673 In *2014 Information Security for South Africa*. IEEE, Johannesburg, South Africa, 1–6. <https://doi.org/10.1109/ISSA.2014.6950490>
- 1674 [103] Nkosinathi Mpofo and Wynand J C van Staden. 2017. Evaluating the severity of trust to identity- management-as-a-service. In *2017 Information
1675 Security for South Africa (ISSA)*. IEEE, Johannesburg, South Africa, 83–89. <https://doi.org/10.1109/ISSA.2017.8251778>
- 1676 [104] Nitin Naik and Paul Jenkins. 2016. A Secure Mobile Cloud Identity: Criteria for Effective Identity and Access Management Standards. In *2016
1677 4th IEEE International Conference on Mobile Cloud Computing, Services, and Engineering (MobileCloud)*. IEEE, Oxford, United Kingdom, 89–90.
1678 <https://doi.org/10.1109/MobileCloud.2016.22>
- 1679 [105] Clifford Neuman, Sam Hartman, Kenneth Raeburn, and Taylor Yu. 2005. The Kerberos Network Authentication Service (V5). RFC 4120.
1680 <https://doi.org/10.17487/RFC4120>
- 1681 [106] James Nicholson, Lynne Coventry, and Pam Briggs. 2013. Age-Related Performance Issues for PIN and Face-Based Authentication Systems. In
1682 *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Paris, France) (CHI '13). Association for Computing Machinery, New
1683 York, NY, USA, 323–332. <https://doi.org/10.1145/2470654.2470701>
- 1684 [107] Martin Nuss, Alexander Puchta, and Michael Kunz. 2018. Towards Blockchain-Based Identity and Access Management for Internet of Things
1685 in Enterprises. In *Trust, Privacy and Security in Digital Business*, Steven Furnell, Haralambos Mouratidis, and Günther Pernul (Eds.). Springer
1686 International Publishing, Cham, 167–181. https://doi.org/10.1007/978-3-319-98385-1_12
- 1687 [108] Raymond J. Offen and Ross Jeffery. 1997. Establishing software measurement programs. *IEEE Software* 14, 2 (Mar 1997), 45–53. <https://doi.org/10.1109/52.582974>
- 1688 [109] Lawrence O’Gorman. 2003. Comparing passwords, tokens, and biometrics for user authentication. *Proc. IEEE* 91, 12 (Dec 2003), 2021–2040.
1689 <https://doi.org/10.1109/JPROC.2003.819611>
- 1690 [110] Wonseok Oh and Alain Pinsonneault. 2007. On the Assessment of the Strategic Value of Information Technologies: Conceptual and Analytical
1691 Approaches. *MIS Quarterly* 31, 2 (2007), 239–265. <https://doi.org/10.2307/25148790>
- 1692 [111] Aleksandr Ometov, Sergey Bezzateev, Niko Mäkitalo, Sergey Andreev, Tommi Mikkonen, and Yevgeni Koucheryavy. 2018. Multi-Factor Authentici-
1693 cation: A Survey. *Cryptography* 2, 1 (2018), 1–31. <https://doi.org/10.3390/cryptography2010001>
- 1694 [112] Ertem Osmanoglu. 2013. *Identity and access management: business performance through connected intelligence*. Elsevier, Waltham, MA, USA.
- 1695 [113] Sven Overhage, Dominik Q. Birkmeier, and Sebastian Schlauderer. 2012. Quality Marks, Metrics, and Measurement Procedures for Business
1696 Process Models. *Business & Information Systems Engineering* 4, 5 (01 Oct 2012), 229–246. <https://doi.org/10.1007/s12599-012-0230-8>
- 1697 [114] Simon Parkinson and Saad Khan. 2022. A Survey on Empirical Security Analysis of Access-Control Systems: A Real-World Perspective. *ACM
1698 Comput. Surv.* 55, 6, Article 123 (dec 2022), 28 pages. <https://doi.org/10.1145/3533703>
- 1699 [115] G. Peterson. 2006. Introduction to identity management risk metrics. *IEEE Security Privacy* 4, 4 (July 2006), 88–91. <https://doi.org/10.1109/MSP.2006.94>
- 1700 [116] Andreas Pfitzmann and Marit Hansen. 2010. A terminology for talking about privacy by data minimization: Anonymity, unlinkability, undetectability,
1701 unobservability, pseudonymity, and identity management.
- 1702 [117] Leo L. Pipino, Yang W. Lee, and Richard Y. Wang. 2002. Data Quality Assessment. *Commun. ACM* 45, 4 (apr 2002), 211–218. <https://doi.org/10.1145/505248.506010>
- 1703 [118] Uthpala Subodhani Premarathne, Ibrahim Khalil, Zahir Tari, and Albert Zomaya. 2017. Cloud-Based Utility Service Framework for Trust Negotiations
1704 Using Federated Identity Management. *IEEE Transactions on Cloud Computing* 5, 2 (April 2017), 290–302. <https://doi.org/10.1109/TCC.2015.2404816>
- 1705 [119] Rosanne Price and Graeme Shanks. 2016. A Semiotic Information Quality Framework: Development and Comparative Analysis. In *Enacting
1706 Research Methods in Information Systems: Volume 3*, Leslie P. Willcocks, Chris Sauer, and Mary C. Lacity (Eds.). Springer International Publishing,
1707 Cham, 219–250. https://doi.org/10.1007/978-3-319-29272-4_7
- 1708 [120] Alexander Puchta, Fabian Böhm, and Günther Pernul. 2019. Contributing to Current Challenges in Identity and Access Management with Visual
1709 Analytics. In *Data and Applications Security and Privacy XXXIII*, Simon N. Foley (Ed.). Springer International Publishing, Cham, 221–239.
- 1710 [121] Alexander Puchta, Sebastian Groll, and Günther Pernul. 2021. Leveraging Dynamic Information for Identity and Access Management: An Extension
1711 of Current Enterprise IAM Architecture. In *Proceedings of the 7th International Conference on Information Systems Security and Privacy - ICISSP*.
INSTICC, SciTePress, Online Streaming, 611–618. <https://doi.org/10.5220/0010315706110618>
- 1712 [122] Nick Ragouzis, John Hughes, Rob Philpott, Eve Maler, Paul Madsen, and Tom Scavo. 2008. *Security Assertion Markup Language (SAML) V2.0 Technical
1713 Overview*. Committee Draft 02 sstc-saml-tech-overview-2.0-cd-02. OASIS. <https://www.oasis-open.org/committees/download.php/27819/sstc-saml-tech-overview-2.0-cd-02.pdf> Committee Specification 02.

- 1717 [123] Erik Rissanen. 2013. *eXtensible Access Control Markup Language (XACML) Version 3.0*. OASIS Standard XACML-V3.0. OASIS. [http://docs.oasis-](http://docs.oasis-open.org/xacml/3.0/xacml-3.0-core-spec-os-en.pdf)
1718 [open.org/xacml/3.0/xacml-3.0-core-spec-os-en.pdf](http://docs.oasis-open.org/xacml/3.0/xacml-3.0-core-spec-os-en.pdf)
- 1719 [124] Pawel Rotter. 2008. A Framework for Assessing RFID System Security and Privacy Risks. *IEEE Pervasive Computing* 7, 2 (April 2008), 70–77.
1720 <https://doi.org/10.1109/MPRV.2008.22>
- 1721 [125] Denis Royer. 2008. Enterprise Identity Management. In *The Future of Identity in the Information Society*, Simone Fischer-Hübner, Penny Duquenoy,
1722 Albin Zuccato, and Leonardo Martucci (Eds.). Springer US, Boston, MA, 433–446.
- 1723 [126] Zhang Rui and Zheng Yan. 2019. A Survey on Biometric Authentication: Toward Secure and Privacy-Preserving Identification. *IEEE Access* 7
1724 (2019), 5994–6009. <https://doi.org/10.1109/ACCESS.2018.2889996>
- 1725 [127] Scott Ruoti and Kent Seamons. 2016. Standard Metrics and Scenarios for Usable Authentication. In *Twelfth Symposium on Usable Privacy and*
1726 *Security (SOUPS 2016)*. USENIX Association, Denver, CO, 1–2. [https://www.usenix.org/conference/soups2016/workshop-program/way2016/](https://www.usenix.org/conference/soups2016/workshop-program/way2016/presentation/ruoti_metrics)
[presentation/ruoti_metrics](https://www.usenix.org/conference/soups2016/workshop-program/way2016/presentation/ruoti_metrics)
- 1727 [128] Riseul Ryu, Soonja Yeom, Soo-Hyung Kim, and David Herbert. 2021. Continuous Multimodal Biometric Authentication Schemes: A Systematic
1728 Review. *IEEE Access* 9 (2021), 34541–34557. <https://doi.org/10.1109/ACCESS.2021.3061589>
- 1729 [129] Nat Sakimura, John Bradley, Michael B. Jones, Breno de Medeiros, and Chuck Mortimore. 2014. OpenID Connect Core 1.0 incorporating errata set
1730 1. https://openid.net/specs/openid-connect-core-1_0.html
- 1731 [130] Mathias Sallé. 2004. IT Service Management and IT Governance: Review, Comparative Analysis and their Impact on Utility Computing.
1732 <https://www.hp.lhp.com/techreports/2004/HPL-2004-98.pdf>
- 1733 [131] Pierangela Samarati and Sabrina Capitani de Vimercati. 2001. Access Control: Policies, Models, and Mechanisms. In *Foundations of Security*
1734 *Analysis and Design*, Riccardo Focardi and Roberto Gorrieri (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 137–196.
- 1735 [132] Rajarajan Sampath and Deepak Goel. 2006. RATING: rigorous assessment of trust in identity management. In *First International Conference on*
1736 *Availability, Reliability and Security (ARES '06)*. IEEE, Vienna, Austria, 10 pp.–23. <https://doi.org/10.1109/ARES.2006.103>
- 1737 [133] Ravi S. Sandhu. 1998. Role-based Access Control. Portions of this chapter have been published earlier in Sandhu et al. (1996), Sandhu (1996),
1738 Sandhu and Bhamidipati (1997), Sandhu et al. (1997) and Sandhu and Feinstein (1994). In *Advances in Computers*, Marvin V. Zelkowitz (Ed.).
1739 *Advances in Computers*, Vol. 46. Elsevier, online, 237–286. [https://doi.org/10.1016/S0065-2458\(08\)60206-5](https://doi.org/10.1016/S0065-2458(08)60206-5)
- 1740 [134] Ravi S. Sandhu and Pierangela Samarati. 1994. Access control: principle and practice. *IEEE communications magazine* 32, 9 (1994), 40–48.
- 1741 [135] Frank Schell, Jochen Dinger, and Hannes Hartenstein. 2009. Performance Evaluation of Identity and Access Management Systems in Federated
1742 Environments. In *Scalable Information Systems*, Peter Mueller, Jian-Nong Cao, and Cho-Li Wang (Eds.). Springer Berlin Heidelberg, Berlin,
1743 Heidelberg, 90–107.
- 1744 [136] Daniel Schlette, Fabian Böhm, Marco Caselli, and Günther Pernul. 2021. Measuring and visualizing cyber threat intelligence quality. *International*
1745 *Journal of Information Security* 20, 1 (01 Feb 2021), 21–38. <https://doi.org/10.1007/s10207-020-00490-y>
- 1746 [137] Roland Schlöglhofer and Johannes Sametinger. 2012. Secure and Usable Authentication on Mobile Devices. In *Proceedings of the 10th International*
1747 *Conference on Advances in Mobile Computing & Multimedia (Bali, Indonesia) (MoMM '12)*. Association for Computing Machinery, New York, NY,
1748 USA, 257–262. <https://doi.org/10.1145/2428955.2429004>
- 1749 [138] Andre Schrimpf, Andreas Drechsler, and Konstantinos Dagianis. 2021. Assessing Identity and Access Management Process Maturity: First
1750 Insights from the German Financial Sector. *Information Systems Management* 38, 2 (2021), 94–115. <https://doi.org/10.1080/10580530.2020.1738601>
1751 arXiv:<https://doi.org/10.1080/10580530.2020.1738601>
- 1752 [139] Ahmed Seffah, Mohammad Donyaee, Rex B. Kline, and Harkirat K. Padda. 2006. Usability measurement and metrics: A consolidated model.
1753 *Software Quality Journal* 14, 2 (01 Jun 2006), 159–178. <https://doi.org/10.1007/s11219-006-7600-8>
- 1754 [140] Ahmed Seffah, N. Kececi, and M. Donyaee. 2001. QUIM: a framework for quantifying usability metrics in software quality models. In *Proceedings*
1755 *Second Asia-Pacific Conference on Quality Software*. IEEE, NW Washington, DC, United States, 311–318. <https://doi.org/10.1109/APAQS.2001.990036>
- 1756 [141] Sean M. Segreti, William Melicher, Saranga Komanduri, Darya Melicher, Richard Shay, Blase Ur, Lujo Bauer, Nicolas Christin, Lorrie Faith
1757 Cranor, and Michelle L. Mazurek. 2017. Diversify to Survive: Making Passwords Stronger with Adaptive Policies. In *Thirteenth Symposium on*
1758 *Usable Privacy and Security (SOUPS 2017)*. USENIX Association, Santa Clara, CA, 1–12. [https://www.usenix.org/conference/soups2017/technical-](https://www.usenix.org/conference/soups2017/technical-sessions/presentation/segreti)
1759 [sessions/presentation/segreti](https://www.usenix.org/conference/soups2017/technical-sessions/presentation/segreti)
- 1760 [142] Nirojan Selvanathan, Dileepa Jayakody, and Violeta Damjanovic-Behrendt. 2019. Federated Identity Management and Interoperability for
1761 Heterogeneous Cloud Platform Ecosystems. In *Proceedings of the 14th International Conference on Availability, Reliability and Security* (Canterbury,
1762 CA, United Kingdom) (*ARES '19*). Association for Computing Machinery, New York, NY, USA, Article 103, 7 pages. [https://doi.org/10.1145/3339252.](https://doi.org/10.1145/3339252.3341492)
1763 [3341492](https://doi.org/10.1145/3339252.3341492)
- 1764 [143] Daniel Servos and Sylvia L. Osborn. 2017. Current Research and Open Problems in Attribute-Based Access Control. *ACM Comput. Surv.* 49, 4,
1765 Article 65 (jan 2017), 45 pages. <https://doi.org/10.1145/3007204>
- 1766 [144] Christopher Staite and Rami Bahsoon. 2012. Evaluating Identity Management Architectures. In *Proceedings of the 3rd International ACM SIGSOFT*
1767 *Symposium on Architecting Critical Systems* (Bertinoro, Italy) (*ISARCS '12*). Association for Computing Machinery, New York, NY, USA, 11–20.
1768 <https://doi.org/10.1145/2304656.2304659>
- [145] Diane M. Strong, Yang W. Lee, and Richard Y. Wang. 1997. Data Quality in Context. *Commun. ACM* 40, 5 (may 1997), 103–110. <https://doi.org/10.1145/253769.253804>
- Manuscript submitted to ACM

- 1769 [146] Mark Thomas. 2021. Using ITIL 4 and COBIT 2019 to create an integrated I&T Framework. [https://www.axelos.com/resource-hub/white-](https://www.axelos.com/resource-hub/white-paper/using-til-cobit-2019-create-integrated-environment)
1770 [paper/using-til-cobit-2019-create-integrated-environment](https://www.axelos.com/resource-hub/white-paper/using-til-cobit-2019-create-integrated-environment)
- 1771 [147] Bianca Trinkenreich, Gleison Santos, and Monalessa Perini Barcellos. 2018. SINIS: A QOM+Strategies-based approach for identifying goals,
1772 strategies and indicators for IT services. *Information and Software Technology* 100 (2018), 147–164. <https://doi.org/10.1016/j.infsof.2018.04.006>
- 1773 [148] United States Congress. 1996. Health Insurance Portability and Accountability Act of 1996.
- 1774 [149] United States Congress. 2002. Sarbanes-Oxley Act of 2002. Corporate responsibility.
- 1775 [150] Irene Vanderfeesten, Jorge Cardoso, Jan Mendling, Hajo A. Reijers, and Wil Aalst, van der. 2007. *Quality metrics for business process models*. Future
1776 Strategies. Lighthouse Point, Florida, USA, 179–190.
- 1777 [151] Carlos Villarrubia, Eduardo Fernández-Medina, and Mario Piattini. 2006. Metrics of Password Management Policy. In *Computational Science and Its*
1778 *Applications - ICCSA 2006*, Marina Gavrilova, Osvaldo Gervasi, Vipin Kumar, C. J. Kenneth Tan, David Taniar, Antonio Laganá, Youngsong Mun,
1779 and Hyunseung Choo (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 1013–1023. https://doi.org/10.1007/11751595_106
- 1780 [152] Emanuel von Zezschwitz, Alexander De Luca, and Heinrich Hussmann. 2014. Honey, I Shrunk the Keys: Influences of Mobile Devices on Password
1781 Composition and Authentication Performance. In *Proceedings of the 8th Nordic Conference on Human-Computer Interaction: Fun, Fast, Foundational*
1782 *(Helsinki, Finland) (NordCHI '14)*. Association for Computing Machinery, New York, NY, USA, 461–470. <https://doi.org/10.1145/2639189.2639218>
- 1783 [153] Isabel Wagner and David Eckhoff. 2018. Technical Privacy Metrics: A Systematic Survey. *ACM Comput. Surv.* 51, 3, Article 57 (jun 2018), 38 pages.
1784 <https://doi.org/10.1145/3168389>
- 1785 [154] Yair Wand and Richard Y. Wang. 1996. Anchoring Data Quality Dimensions in Ontological Foundations. *Commun. ACM* 39, 11 (nov 1996), 86–95.
1786 <https://doi.org/10.1145/240455.240479>
- 1787 [155] Richard Y. Wang, Veda C. Storey, and Christopher P. Firth. 1995. A framework for analysis of data quality research. *IEEE transactions on knowledge*
1788 *and data engineering* 7, 4 (1995), 623–640.
- 1789 [156] Richard Y. Wang and Diane M. Strong. 1996. Beyond Accuracy: What Data Quality Means to Data Consumers. *J. Manage. Inf. Syst.* 12, 4 (mar
1790 1996), 5–33. <https://doi.org/10.1080/07421222.1996.11518099>
- 1791 [157] Jorge Werner, Carla Merkle Westphall, and Carlos Becker Westphall. 2017. Cloud identity management: A survey on privacy strategies. *Computer*
1792 *Networks* 122 (2017), 29–42. <https://doi.org/10.1016/j.comnet.2017.04.030>
- 1793 [158] Alan F Westin. 1968. Privacy and freedom. *Washington and Lee Law Review* 25, 1 (1968), 166.
- 1794 [159] Phillip J Windley. 2005. *Digital Identity: Unmasking identity management architecture (IMA)*. O'Reilly Media, Inc., Sebastopol, CA.
- 1795 [160] Peter Wood. 2005. Implementing identity management security - an ethical hacker's view. *Network Security* 2005, 9 (2005), 12–15. [https://doi.org/10.1016/S1353-4858\(05\)70282-8](https://doi.org/10.1016/S1353-4858(05)70282-8)
- 1796 [161] Chengcheng Xiang, Yudong Wu, Bingyu Shen, Mingyao Shen, Haochen Huang, Tianyin Xu, Yuanyuan Zhou, Cindy Moore, Xinxin Jin, and
1797 Tianwei Sheng. 2019. Towards Continuous Access Control Validation and Forensics. In *Proceedings of the 2019 ACM SIGSAC Conference on*
1798 *Computer and Communications Security (London, United Kingdom) (CCS '19)*. Association for Computing Machinery, New York, NY, USA, 113–129.
1799 <https://doi.org/10.1145/3319535.3363191>
- 1800 [162] Dianxiang Xu, Lijo Thomas, Michael Kent, Tejeddine Mouelhi, and Yves Le Traon. 2012. A Model-Based Approach to Automated Testing of Access
1801 Control Policies. In *Proceedings of the 17th ACM Symposium on Access Control Models and Technologies (Newark, New Jersey, USA) (SACMAT '12)*.
1802 Association for Computing Machinery, New York, NY, USA, 209–218. <https://doi.org/10.1145/2295136.2295173>
- 1803 [163] Zhongyuan Xu. 2014. *Mining Meaningful Role-Based and Attribute-Based Access Control Policies*. Ph. D. Dissertation. State University of New York
1804 at Stony Brook.
- 1805 [164] Yi Yang, Kheng Cher Yeo, Sami Azam, Asif Karim, Ronju Ahammad, and Rakib Mahmud. 2020. Empirical Study of Password Strength Meter
1806 Design. In *2020 5th International Conference on Communication and Electronics Systems (ICCES)*. IEEE, Coimbatore, India, 436–442. <https://doi.org/10.1109/ICCES48766.2020.9137964>
- 1807 [165] Emrah Yasasin and Guido Schryen. 2015. Requirements for IT Security Metrics – An Argumentation Theory Based Approach. In *23rd European*
1808 *Conference on Information Systems (ECIS 2015)*. AIS, Münster, Germany, 1–16. <https://doi.org/10.18151/7217537>
- 1809 [166] Kurt Zeilenga. 2006. Lightweight Directory Access Protocol (LDAP): Directory Information Models. RFC 4512. <https://doi.org/10.17487/RFC4512>
- 1810 [167] Xiaotong Zhou, Debiao He, Jianting Ning, Min Luo, and Xinyi Huang. 2023. AADEC: Anonymous and Auditable Distributed Access Control for
1811 Edge Computing Services. *IEEE Transactions on Information Forensics and Security* 18 (2023), 290–303. <https://doi.org/10.1109/TIFS.2022.3220030>

1812 Received 01 January 2023; revised 01 January 2023; accepted 01 January 2023

1813

1814

1815

1816

1817

1818

1819

1820

Manuscript submitted to ACM

3 Transaction Logs in Access Control: Leveraging an Under-Utilized Data Source

Current status:	Published
Conference:	39th Annual IFIP WG 11.3 Conference on Data and Applications Security and Privacy (DBSec 2025), Gjøvik, June 23-25, 2025
Date of acceptance:	May 01, 2025
Full citation:	Sascha Kern, Thomas Baumer, Raphael Neudert, and Günther Pernul. Transaction Logs in Access Control: Leveraging an Under-Utilized Data Source. In <i>IFIP Annual Conference on Data and Applications Security and Privacy</i> , pages 413-424. Springer, 2025.
Authors contributions:	Sascha Kern 40% Thomas Baumer 25% Raphael Neudert 25% Günther Pernul 10%

Conference Description: The 39th edition of the Annual IFIP WG 11.3 Conference on Data and Applications Security and Privacy (DBSec 2025) will take place in Gjøvik, Norway. The conference brings together researchers, practitioners, and experts from academia, industry, and government to share their cutting-edge findings and insights in all theoretical and practical aspects of data protection, privacy, and applications security.



Transaction Logs in Access Control: Leveraging an Under-Utilized Data Source

Sascha Kern¹, Thomas Baumer¹ (✉), Raphael Neudert¹,
and Günther Pernul²

¹ Nexis GmbH, Rudolf-Vogt-Straße 6, 93053 Regensburg, Germany
{sascha.kern,thomas.baumer,raphael.neudert}@nexis-secure.com

² University of Regensburg, Universitätsstraße 31, 93053 Regensburg, Germany
guenther.pernul@ur.de

Abstract. Maintaining access control policies is an ongoing process to ensure required but not excessive authorizations. Organizations thus leverage various data sources to ease this maintenance. Among these data sources are access control matrices, attributes, access logs, and transaction logs. While research reasonably covers the former data sources, the potential of transaction logs remains untapped. We pave the way for transaction logs as a data source in access control by (i) expressing them with a formalization, (ii) pinpointing them in typical Identity and Access Management (IAM) infrastructures, and (iii) grounding them in IAM processes. We conclude that access control transaction logs are valuable data sources for improving analytical capabilities for IAM.

Keywords: Access Control · Transaction Logs · Data Quality

1 Introduction

Existing research has devoted considerable attention to the initial creation of high-quality access control policies [12, 23]. Once created, policies lose quality due to changing environmental conditions, incorrect or suboptimal policy adjustments, or over-granting [29]. Access control policies must thus be maintained in an ongoing manner to remain accurate. This is intuitive for static access control models like Role-Based Access Control (RBAC). Dynamic models, like e.g. Attribute-Based Access Control (ABAC), which grant access based on attribute values, cannot eliminate this problem either, since their policies also reflect the organizational context at the time of creation and can become outdated due to changes in the environment (e.g., new regulatory requirements, organizational structures or application systems). In practice, it is prevalent to analyze access control policies manually to identify inaccurate authorizations, e.g., by Access Reviews [18]. The effectiveness of manual maintenance is limited, as authorization data is often vast and challenging to analyze [20].

Approaches that aim to automate parts of policy maintenance exist. Still, they are limited by the availability and quality of required data [21]: The most

414 S. Kern et al.

important and readily available data types are the access control matrix and attributes. The access control matrix and its variants describe the effective authorizations defined by the access control policies. At the same time, attributes offer valuable context information that helps understand policies' semantics and identify patterns and outliers. However, these data only describe the current state of access control, which is likely outdated and thus equally erroneous, as in the policies. Therefore, policy maintenance with no further data reproduces existing errors and is limited to environments with high data quality. Another data source used for policy maintenance is access logs [34]. They allow for identifying inaccurate authorizations by analyzing which authorizations have been invoked in a certain period. However, the availability of access logs can be limited, and their use requires a certain level of interpretation.

We introduce transaction logs as a data source for maintaining access control policies. Transaction logs are structured change histories for relevant IAM data. We provide a formal definition and show that transaction logs are available in IAM infrastructures. We ground the validity of access control transaction logs in IAM processes. Finally, we conduct a case study with a real-world enterprise to verify our findings. In summary, the contributions of this work are as follows:

- C1** *We formalize transaction logs in access control.*
- C2** *We locate transaction logs in typical, centralized IAM infrastructures.*
- C3** *We ground the analytical value of transaction logs in IAM processes.*

This work is structured as follows. We position our work with related work in Sect. 2. Section 3 analyzes the utilization of transaction logs in three steps: (i) We formalize transaction logs in IAM. (ii) We analyze how to locate transaction logs in IAM infrastructures. (iii) We ground the validity of transaction logs on typical IAM processes. Section 4 concludes this work.

2 Related Work

Research proposes methods to improve the quality of access control policies. They differ in terms of the addressed quality dimensions and required data. We present relevant approaches along with their respective data reliance and known limitations. Table 1 summarizes the known data usages and limitations.

The *access control matrix* displays the authorizations granted by a policy set as they are. Theoretically, it can identify patterns in existing authorizations and adjust them to remove outliers. For example, one identity sharing all authorizations except one with ten other identities could be assigned the missing authorization. However, without further data to provide context information, the analytical power of the access control matrix is limited. Some authors propose visualizations based on the access control matrix to improve the data understanding of policy engineers [22, 25]. Other works that process the access control matrix to improve policy quality typically use it to recreate existing authorizations, while modifying policies to improve other quality dimensions such as complexity [6, 33], redundancy, or the number of conflicts [26].

Table 1. Data types and their known uses and limitations for access control policy maintenance.

Data Category	Known Use	Known Limitations
Access Control Matrix	Recreating authorizations	Analytical power
Attributes	Identifying authorization patterns and outliers	Detectable inaccuracies Attribute and ACP quality
Access Logs	Estimating inaccurate authorizations with access invocations	Availability Sensitive data Detectable inaccuracies
Transaction Logs	Draw authorization insights from IAM processes	No formalization or localization Missing grounding

Attributes provide valuable information about real-world meanings of access control policies. They are commonly used by policy mining approaches to create semantically meaningful policies. We argue ABAC mining approaches might serve as blueprints for attribute-based policy maintenance. Note that using ABAC policies alone does not fully eliminate the problem of policy maintenance, since the attribute rules can outdate just like other access control models. Suppose no data other than the present state attributes and the access control matrix are available. In that case, policy maintenance is limited to outlier detection, and only a relatively small fraction of authorizations can be identified as inaccurate. Moreover, outlier detection approaches are limited by their strong reliance on attribute quality and authorization accuracy. Thus, a policy set cannot be maintained if its quality is too low.

Several approaches use *access logs* to identify excessive authorizations [14, 30], missing authorizations [4], or both [5, 19]. Access logs are records of historic permission invocations in the form $\langle S, O, A, R \rangle$, with S being the subject that requests access, O and A being the requested object and action, and R being the response returned by the access control mechanism, i.e., a *permit* or *deny* decision. The shared concept for identifying excessive authorizations with access logs is to search for authorizations not used in a certain period (e.g., last year). Authorizations that are not needed are excessive and can thus be removed according to the principle of least privilege. However, if a user holds an excessive authorization and uses it, it cannot be detected using access logs. Missing authorizations are identified by analyzing authorization requests that were denied. However, this method requires further interpretation since the mere request of an authorization does not necessarily mean that an identity *should* have it. Regarding availability, core applications like operating and database management systems often generate detailed access logs. In contrast, higher-level applications tend to create fewer, as it is not a prioritized use case. Another limitation can stem from legal requirements, as access logs enable user monitoring, which is restricted by some privacy protection laws. Access logs can be a valuable data source for identify-

416 S. Kern et al.

ing inaccurate authorizations. However, their utilization can be limited by their availability, the sensitivity of their content, and the amount of inaccuracies they can reveal. Some approaches like [15,30] also use access logs to identify behavior anomalies, e.g., many access requests that occurred within a short period, and revoke the related authorizations. However, these authorizations are not revoked for being excessive, but as an immediate threat response.

Transaction logs are structured records of data changes. Changes to access control policies and digital identities are particularly relevant in access control. Transaction logs allow for the tracking of changes that have occurred and the restoration of previous data states. Despite inconsistent terminology, they are a well-established data type used in other application fields of IT, e.g., database systems [9], software version control systems [7], or digital ledger technologies [3]. Some authors have emphasized the analytical value of (structured) data change logs in IAM: Molloy et al. [24] propose to identify job change events, recognizable by the pattern of an identity losing and receiving several authorizations within a short time frame, to determine semantically meaningful permission groupings for roles. They also argue that historical update information can be used to identify legacy permissions that should no longer be authorized altogether. Mitra et al. [23] propose to use transaction logs to identify outdated roles that require maintenance. Both Strembeck [31] and Fuchs et al. [10] name trace management as requirements for role model maintenance that provide role engineers with valuable context information. Hein et al. [16] use transaction logs to enable administrators to rollback suboptimal changes to a role set. Hunt et al. [17] also suggest standardizing transaction logs for the provisioning standard System for Cross-domain Identity Management (SCIM) and name replication of IAM systems as a possible use case. However, none of these works introduces a formalization or analyzes the infrastructure localization of transaction logs or their utilization for policy maintenance.

This work offers a foundation for analyzing authorization structures with transaction logs. Established IAM processes provide a structure for the analysis. To the best of our knowledge, this is the first work to formally define transaction logs in the context of access control.

3 Transaction Logs in Identity and Access Management

This section examines how transaction logs can be used to provide analysis support for the maintenance of access control policies. Section 3.1 formally defines transaction logs. Section 3.2 examines where transaction logs can be collected from a schematic IAM infrastructure. Section 3.3 grounds access control transaction logs in typical IAM processes.

3.1 Formalization of Transaction Logs

For a precise and common notation throughout our work, we formalize transaction logs, including their preliminaries, the transactions themselves, and their

interactions. We show that transaction logs for access control do not differ substantially from related concepts for database systems or version control. Figure 1 depicts the basic idea of transactions and their interactions, which we detail in the following paragraphs.




Fig. 1. Transactions and their interactions.

Preliminaries

Let E be the set of all entities, A the set of all attributes assignable to entities, and V the set of all attribute values.

- $e \in E$ is an entity.
- $a \in A$ is an attribute.
- $v \in V$ is an attribute value.

Entities have an arbitrary set of attributes assigned, and an attribute can have an arbitrary amount of values. For access control, entities refer to digital identities or access control policies.

States

- An entity e has an arbitrary amount of states $s^e \in S^e$. $S^e = \{\emptyset, s_1^e, \dots, s_n^e\}$ describes a sequence of $n + 1$ states of an entity $e \in E$, with \emptyset being the empty state. The states are ordered by their occurrence in time. We follow that s refers to an entity $e \in E$ at a specific time.

Transactions

- $TR^e = \{tr_{create}^e, tr_1^e, \dots, tr_{n-1}^e, tr_{delete}^e\}$ is the sequence of transactions for an entity e . It contains the descriptions of state changes, ordered by their occurrence over time.
- We define a transaction $tr^e \in TR^e$ with the tuple $\langle a, v, t \rangle$. A transaction tr_i^e describes a single state change $s_i^e \rightarrow s_{i+1}^e$. Along with a and v , it contains a timestamp t that allows ordering the transactions in TR^e .
- We define a *create* transaction $tr_{create}^e = \langle \emptyset, \emptyset, t \rangle$. It describes the initial creation of e as it is added to E .
- We define a *delete* transaction $tr_{delete}^e = \langle \emptyset, \emptyset, t \rangle$. It describes the removal of e from E .

While transactions describe state changes of entities, interactions use transactions to transition between states.

418 S. Kern et al.

Interactions

- We define an update function $update(s_i^e, tr_i^e) = s_{i+1}^e$. Thus, by applying the next transaction tr_i^e to a given data state s_i^e , the data state is transitioned forward ($s_i^e \rightarrow s_{i+1}^e$).
- $update(\emptyset, tr_{create}^e)$ creates the first non-empty state s_1^e .
- $update(s_n^e, tr_{delete}^e)$ creates the empty state \emptyset .
- We define a rollback function $rollback(s_i^e, tr_{i-1}^e) = s_{i-1}^e$. Thus, by applying the previous transaction tr_{i-1}^e in a rollback function, a state is transitioned backward ($s_i^e \rightarrow s_{i-1}^e$).
- $rollback(\emptyset, tr_{delete}^e)$ creates the last non-empty state s_n^e .
- $rollback(s_1^e, tr_{create}^e)$ creates the empty state \emptyset .

Transaction thus enable two interactions with entity states: update and rollback. While the update function proceeds the entity state forward in time, the rollback function reverses the entity to a previous state. Both functions thus compute previous or new transactions to reach the desired entity states.

Transaction Log

- We define a transaction log $L \supset TR^e, \forall e \in E$, i.e. the sequence of transactions observed in a given time span for all observed entities. A transaction log entry is thus a single transaction for a specific entity.

Finally, transaction logs are the main focus of our study. Along with the current data states for these entities, transaction logs allow to analyze data changes, and to restore previous data states. Since this depiction of transaction logs is not specific to access control, we consider access control transaction logs a subset of transaction logs in general computer science.

3.2 Localization in IAM Infrastructures

We analyze the availability of transaction logs in centralized IAM infrastructures. We introduce a schematic architecture of centralized IAM and analyze transaction logs' sources and collection constraints.

Schematic IAM Infrastructure. An IAM infrastructure controls user authentication and authorization. Due to the limited scope of this work, we do not consider the authentication details and focus on centralized IAM infrastructures typical for mid to large-sized organizations [8]. To execute user authorization, an IAM infrastructure must store user identities, policies, and supplemental data (e.g., department structures). While this data is not limited to a specific type of storage system, it is often found in relational databases, HR systems, directory systems, or meta-directories [1, 11]. Any application system that relies on authorization functionality is a target for access control and is thus called a *target system*. An IAM infrastructure can include numerous target systems such as workstations or business application systems (e.g., a banking

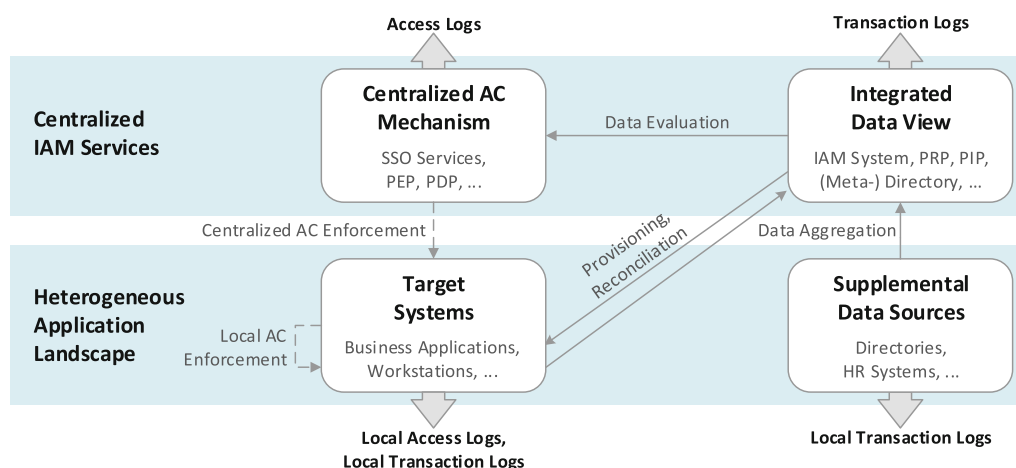


Fig. 2. Schematic centralized [8] IAM infrastructure and possible log data generation.

system). The mechanism that evaluates policies to enforce the defined authorizations is called an *access control mechanism* [27]. It can either reside locally within the target system (which is the isolated model [8]) or exist as a central service in the IAM infrastructure, as defined e.g. by the eXtensible Access Control Markup Language (XACML) standard, Single Sign On (SSO) standards like Open Authorization (OAuth) or the Security Assertion Markup Language (SAML), or provision standards like SCIM or Service Provisioning Markup Language (SPML).

To enable modeling and managing access control policies throughout an organization, an IAM infrastructure must aggregate data from the numerous target systems and supplemental data sources in an integrated data view [1, 11]. This data view includes the authorizations, user identities (and possibly additional information), and the policies defining the resulting authorizations. Common examples of systems providing centralized IAM data views are industrial IAM systems, the Policy Retrieval Point (PRP) and Policy Information Point (PIP) of the XACML standard, (meta-) directories, or relational database views. For decentralized access control mechanisms to work, newly created or changed user identities and access control policies must be provisioned into the target systems in a *provisioning process*, and data inconsistencies must be identified and rectified in a *reconciliation process*. These data aggregation and synchronization tasks are commonly executed by IAM systems (like SailPoint IIQ, One Identity Manager, Oracle IAM, or Microsoft Entra), which also offer extensive logging and data overview functionality. A centralized access control mechanism can reduce or eliminate the need for data provisioning and reconciliation but requires high system interoperability and data standardization.

In summary, the schematic IAM services operate on two infrastructure layers: The heterogeneous application landscape includes numerous target systems and supplemental data sources. On top of it resides a centralized IAM service layer, including an integrated IAM data view and possibly a centralized access control

420 S. Kern et al.

mechanism. Any application may generate various types of log data, e.g. error records or debug information. The generation of transaction logs is not as arbitrary and will be examined in the following. Figure 2 summarizes the schematic IAM infrastructure and the types of log data generated by its components.

Localization of Transaction Logs. As defined in Sect. 3.1, transaction logs are structured records of changes to IAM data. The relevant entities include digital identities, access control policies, and possibly additional information. The integrated data view offers an organization-wide summary of this data. A full transaction log can thus be generated by simply monitoring changes in the integrated data view and logging them in a structured way. This is trivial since it only requires logging all observed create, update, and delete operations on the relevant data, a standard functionality in common database management systems. Industrial IAM systems also need to generate detailed change logs by default to comply with legal regulations¹ which require organizations to track when and how employees were granted rights and how access control rules were modified. Since an integrated IAM data view is a prerequisite for the organization-wide modeling and management of access control policies, we argue that transaction logs are easy to obtain. On principle, these transactions can also be logged locally throughout the heterogeneous application landscape. However, utilizing local transaction logs requires log collection and normalization effort.

3.3 Grounding Transaction Logs on IAM Processes

Transaction log entries provide analytical value because they document events that can be mapped to meaningful real-world activities. IAM Processes provide a structure for such activities to be located and analyzed. A process instance (i.e., a single process execution) spans over a certain execution time and covers a set of logged events that document its execution. Mapping logged events to a process is not always trivial: It requires identifying the process from the log events and then determining which events were part of which process instance. The process mining research domain provides valuable tools for this mapping [32]. For the maintenance of access control policies, IAM processes that change existing authorizations or the composition of policies are especially interesting. Several of these processes are standardized to a certain degree and documented in frameworks such as identity life cycle models, policy life cycle models, or capability maturity models [28]. This section analyzes three exemplary processes and their possible mapping to transaction logs: Joiner, Mover, and Access Review. We selected them because of their wide adoption and clear grounding for access control transaction logs.

¹ Global, national, and federal regulation efforts on IAM are numerous and heterogeneous in detail. Notable examples include audited compliance or certifications, like the Sarbanes-Oxley Act (SOX), the Basel Accords, the European General Data Privacy Regulation (GDPR), or the ISO 27000 standards.

The Joiner Process is executed when a person enters an organization. It includes the creation of a digital identity with its respective attributes, as well as the initial assignment of the permissions that the person requires. Typical events documenting the creation of a digital identity are new employee records in an HR system or user accounts in numerous target systems. In these examples, new authorizations can be assigned by roles, directory groups, or target system-specific permissions. Identified instances of the Joiner process can reveal a lot of information as they show a complete digital identity that receives all required authorizations shortly after its creation.

Over-granting causes digital identities to accumulate authorizations. As a result, users typically have significantly more authorizations than they should have according to the principle of least privilege. According to an estimate by practitioners, the proportion of excessive authorization commonly amounts to more than 20% of total authorizations [2]. Maintenance approaches that rely on an organization's current attributes and access control matrix are prone to reproducing these errors [21]. The authorizations granted in the Joiner process are unlikely to contain many excessive authorizations since no accumulation has yet taken place. The initially set attributes are also guaranteed to be timely and likely of high data quality. Therefore, the Joiner Process can provide a better data foundation for policy maintenance than an organization's current attributes and access control matrix. At the same time, the Joiner process provides a complete picture of the digital identities and authorizations at the time of their creation, meaning that it allows the identification of complex authorization patterns.

The Mover Process is executed when a person changes their affiliation within the organization, e.g., the department, cost center, or job title. Due to changes in the person's responsibilities, a Mover process execution is often accompanied by changes in their authorizations. A Mover process execution shows in the transaction logs through attribute changes, e.g., a new value for the attribute department, or through a digital identity losing and receiving multiple authorizations in a short time span. The Mover process allows the identification of strong correlations between attributes and authorizations without requiring knowledge of the full authorization structure: If changes of authorizations or attributes frequently occur together, this indicates that they depend on one another. However, identifying complex attribute patterns can be challenging since a Mover process does not necessarily need to change all attributes required for an authorization. For example, employees might require the permission *Network Access for External IT Employees* after moving from marketing to the IT department. If employees were already externals, the Mover process would only show the attribute change *department = IT*, but not *employee type = external*. Consequently, it is easier to detect excessive authorizations with Mover processes than missing ones, since a digital identity missing either one of these attributes would imply that it should not inherit this permission. In contrast, detecting missing permissions requires knowing the complete pattern.

The Access Review Process aims to identify excessive authorizations (and possibly other errors affecting access control policies, e.g. inaccurate attributes).

422 S. Kern et al.

It is periodically carried out by employees who check the authorizations of subjects in their responsibility (e.g. department heads). If an excessive authorization is found, it is revoked [18]. Access Review executions are thus shown in transaction logs through their authorization revocation events. They are a valuable source for analysis because they show confirmed errors in the access control matrix. Like access logs, these identified inaccuracies provide a ground truth that can be used to search for further errors. Groll et al. provide formalizations for the analysis of Access Review results. They propose an approach that identifies further excessive authorizations based on past Access Review decisions [13].

Further Implications. Due to over-granting, access control transaction logs are likely to contain more entries that document users receiving authorization than authorization revocations. Furthermore, frequent over-granting causes authorization grant entries to be inaccurate, more likely than revocation entries. This means that revocation entries are potentially more valuable for analysis than grants but also scarcer. Revocation entries might thus be analyzed with higher priority, provided they occur frequently enough to be significant.

4 Conclusion

Our work contributes toward utilizing transaction logs in access control. We formalize the notion of transaction logs in access control. We depict typical IAM infrastructures and locate access control transaction logs within them. Finally, we ground the evaluation of access control transaction logs in typical IAM processes.

We invite fellow researchers to contribute to access control transaction logs in future work: (i) Association rules based on transaction logs may help to find new rules for automated role assignment, ABAC policies, or errors in existing policies and attributes. (ii) Understanding access control policies is a challenge for their maintainers [20]. Automatically generating descriptions of policies based access control transaction logs might add comprehensibility by adding context, e.g., “This role is typically granted to recently promoted managers”. (iii) A seamless record of transaction logs allows for inspecting past states. Browsing these past states can provide further insights or visualize the development of metrics not yet defined in the past. Overall, we conclude that transaction logs are a promising yet underutilized data source for improving access control.

Acknowledgments. The German Federal Ministry of Education and Research supported the research leading to these results as part of the DEVISE project.

References

1. Baumer, T., Müller, M., Pernul, G.: System for cross-domain identity management (SCIM): survey and enhancement with RBAC. *IEEE Access* **11**, 86872–86894 (2023)

2. Baumer, T., Reitinger, T., Kern, S., Pernul, G.: Digital nudges for access reviews: guiding deciders to revoke excessive authorizations. In: Proceedings of the Twentieth USENIX Conference on Usable Privacy and Security, SOUPS 2024. USENIX Association, USA (2024). <https://doi.org/10.5555/3696899.3696912>
3. Beck, R., Czepluch, J.S., Lollike, N., Malone, S.: Blockchain—the gateway to trust-free cryptographic transactions. In: Twenty-Fourth European Conference on Information Systems (ECIS), İstanbul, Turkey, pp. 1–14. Springer (2016)
4. Benedetti, M., Mori, M.: Parametric RBAC maintenance via max-SAT. In: Proceedings of the 23rd ACM on Symposium on Access Control Models and Technologies, pp. 15–25 (2018)
5. Benedetti, M., Mori, M.: On the use of max-SAT and PDDL in RBAC maintenance. *Cybersecurity* **2**, 1–25 (2019)
6. Benkaouz, Y., Erradi, M., Freisleben, B.: Work in progress: K-nearest neighbors techniques for ABAC policies clustering. In: Proceedings of the 2016 ACM International Workshop on Attribute Based Access Control, pp. 72–75 (2016)
7. Buffardi, K.: Assessing individual contributions to software engineering projects with git logs and user stories. In: Proceedings of the 51st ACM Technical Symposium on Computer Science Education, pp. 650–656 (2020)
8. Cao, Y., Yang, L.: A survey of identity management technology. In: 2010 IEEE International Conference on Information Theory and Information Security, pp. 287–293 (2010). <https://doi.org/10.1109/ICITIS.2010.5689468>
9. Davis, T., Shaw, G., Delaney, K.: SQL Server Transaction Log Management. Simple Talk Pub. (2012)
10. Fuchs, L., Kunz, M., Pernul, G.: Role model optimization for secure role-based identity management. In: European Conference on Information Systems (ECIS), pp. 1–15. AIS, Tel Aviv (2014). <https://epub.uni-regensburg.de/30394/>
11. Fuchs, L., Pernul, G.: Supporting compliant and secure user handling - a structured approach for in-house identity management. In: The Second International Conference on Availability, Reliability and Security (ARES 2007), pp. 374–384. IEEE, Vienna (2007). <https://doi.org/10.1109/ARES.2007.145>
12. Fuchs, L., Pernul, G.: Hydro-hybrid development of roles. In: Information Systems Security: 4th International Conference, ICISS 2008, Hyderabad, India, 16–20 December 2008. Proceedings 4, pp. 287–302. Springer (2008)
13. Groll, S., Kern, S., Fuchs, L., Pernul, G.: Monitoring access reviews by crowd labelling. In: Trust, Privacy and Security in Digital Business: 18th International Conference, TrustBus 2021, Virtual Event, 27–30 September 2021, Proceedings 18, pp. 3–17. Springer (2021)
14. Gunter, C.A., Liebovitz, D., Malin, B.: Experience-based access management: a life-cycle framework for identity and access management systems. *IEEE Secur. Priv.* **9**(5), 48 (2011)
15. Hasel Mehri, G., Wester, I.L., Paci, F., Zannone, N.: Mitigating privilege misuse in access control through anomaly detection. In: Proceedings of the 18th International Conference on Availability, Reliability and Security, ARES 2023. Association for Computing Machinery, New York (2023)
16. Hein, P., Biswas, D., Martucci, L.A., Muhlhauser, M.: Conflict detection and life-cycle management for access control in publish/subscribe systems. In: 2011 IEEE 13th International Symposium on High-Assurance Systems Engineering, pp. 104–111. IEEE (2011)
17. Hunt, P., Cam-Winget, N., Kiser, M., Schreiber, J.: SCIM profile for security event tokens. internet-draft draft-ietf-scim-events-07. Internet Engineering Task Force (2024). Work in Progress

424 S. Kern et al.




18. Jaferian, P., Rashtian, H., Beznosov, K.: To authorize or not authorize: helping users review access policies in organizations. In: 10th Symposium On Usable Privacy and Security (SOUPS 2014), pp. 301–320 (2014)
19. Karimi, L., Aldairi, M., Joshi, J., Abdelhakim, M.: An automatic attribute-based access control policy extraction from access logs. *IEEE Trans. Dependable Secure Comput.* **19**(4), 2304–2317 (2021)
20. Kern, S., Baumer, T., Fuchs, L., Pernul, G.: Maintain high-quality access control policies: an academic and practice-driven approach. In: Atluri, V., Ferrara, A.L. (eds.) *DBSec 2023*. LNCS, vol. 13942, pp. 223–242. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-37586-6_14
21. Kern, S., Baumer, T., Groll, S., Fuchs, L., Pernul, G.: Optimization of access control policies. *J. Inf. Secur. Appl.* **70**, 103301 (2022)
22. Leitner, M., Rinderle-Ma, S.: Anomaly detection and visualization in generative RBAC models. In: *Proceedings of the 19th ACM Symposium on Access Control Models and Technologies*, pp. 41–52 (2014)
23. Mitra, B., Sural, S., Vaidya, J., Atluri, V.: A survey of role mining. *ACM Comput. Surv.* **48**(4) (2016). <https://doi.org/10.1145/2871148>
24. Molloy, I., Chen, H., Li, T., Wang, Q., Li, N., Bertino, E., Calo, S., Lobo, J.: Mining roles with multiple objectives. *ACM Trans. Inf. Syst. Secur. (TISSEC)* **13**(4), 1–35 (2010)
25. Morisset, C., Sanchez, D.: VisABAC: a tool for visualising ABAC policies. In: *ICISSP*, pp. 117–126 (2018)
26. Pang, C., Hansen, D., Maeder, A.: Managing RBAC states with transitive relations. In: *Proceedings of the 2nd ACM Symposium on Information, Computer and Communications Security*, pp. 139–148 (2007)
27. Parkinson, S., Khan, S.: A survey on empirical security analysis of access-control systems: a real-world perspective. *ACM Comput. Surv.* **55**(6) (2022)
28. Schrimpf, A., Drechsler, A., Dagianis, K.: Assessing identity and access management process maturity: first insights from the German financial sector. *Inf. Syst. Manag.* **38**(2), 94–115 (2021)
29. Shen, B., Shan, T., Zhou, Y.: Improving logging to reduce permission over-granting mistakes. In: *32nd USENIX Security Symposium (USENIX Security 2023)*, pp. 409–426. USENIX Association, Anaheim (2023)
30. Skopik, F., Wurzenberger, M., Höld, G., Landauer, M., Kuhn, W.: Behavior-based anomaly detection in log data of physical access control systems. *IEEE Trans. Dependable Secure Comput.* **20**(4), 3158–3175 (2023)
31. Strembeck, M.: Scenario-driven role engineering. *IEEE Secur. Priv.* **8**(1), 28–35 (2010)
32. Van Der Aalst, W.: Process mining: overview and opportunities. *ACM Trans. Manage. Inf. Syst. (TMIS)* **3**(2), 1–17 (2012)
33. Xia, H., Dawande, M., Mookerjee, V.: Role refinement in access control: model and analysis. *INFORMS J. Comput.* **26**(4), 866–884 (2014)
34. Xiang, C., et al.: Towards continuous access control validation and forensics. In: *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, pp. 113–129 (2019)

4 Maintain High-Quality Access Control Policies: An Academic and Practice-Driven Approach.

Current status:	Published
Conference:	37th Annual IFIP WG 11.3 Conference on Data and Applications Security and Privacy (DBSec 2023), Sophia Antipolis, July 19th - July 21th, 2023
Date of acceptance:	May 24, 2023
Full citation:	Sascha Kern, Thomas Baumer, Ludwig Fuchs, and Günther Pernul. Maintain high-quality access control policies: an academic and practice-driven approach. In <i>IFIP Annual Conference on Data and Applications Security and Privacy</i> , pages 223–242. Cham: Springer Nature Switzerland, 2023.
Authors contributions:	Sascha Kern 50% Thomas Baumer 35% Ludwig Fuchs 5% Günther Pernul 10%

Conference Description: The 37th edition of the Annual IFIP WG 11.3 Conference on Data and Applications Security and Privacy (DBSec 2023) will take place in Sophia Antipolis, France. The conference brings together researchers, practitioners, and experts from academia, industry, and government to share their cutting-edge findings and insights in all theoretical and practical aspects of data protection, privacy, and applications security.

Maintain High-Quality Access Control Policies: An Academic and Practice-Driven Approach*

Sascha Kern¹, Thomas Baumer¹, Ludwig Fuchs¹, and Günther Pernul²

¹ Nexis GmbH, Franz-Mayer-Straße 1, Regensburg, 93053, Bavaria, Germany
<https://nexis-secure.com/>

² University of Regensburg, Universitätsstraße 31, Regensburg, 93053, Bavaria,
Germany www.ur.de/informatik-data-science/wi-pernul/startseite

Abstract. Organizations encounter great difficulties in maintaining high-quality Access Control Policies (ACPs). Policies originally modeled and implemented with good quality deteriorate over time, leading to inaccurate authorization decisions and reduced policy maintainability. As a result, security risks arise, delays prevent users from carrying out tasks, and ACP management becomes more expensive and error-prone. In contrast to the initial modeling of ACPs, their long-term maintenance has been addressed scarcely by existing research. This work addresses this research gap with three contributions: First, we provide a detailed problem analysis based on a literature survey and six real-world practitioner expert interviews. Second, we propose a framework that supports organizations in implementing and performing ACP maintenance. Third, we present a maintenance case study in which we implemented maintenance capabilities for a real-world ACP dataset that allowed us to significantly improve its quality.

Keywords: Identity management · Access control · Access control policies · Data quality · Policy maintenance · Security management

1 Introduction

Authorizing users' access to protected resources is a cornerstone of every modern IT security framework. While technologies to enforce well-defined authorizations exist, organizations still struggle with their management: Numerous scientific studies and industry reports highlight major difficulties in adhering to the Principle of Least Privilege (PoLP) [46,34] and point out the high frequency of related IT security vulnerabilities, such as attacks through malicious insiders or hijacking of privileged identities [1,9]. The basis for the definition of IT authorizations are Access Control Policies (ACPs). These machine-processable rules define the user's access to resources. The high-quality modeling of new ACPs has received significant interest in research realms such as policy mining and policy engineering. However, policy modeling is not a one-off effort: Changes within

* The research leading to these results was supported by the German Federal Ministry of Education and Research as part of the DEVISE project (<https://devise.ur.de>).

2 S. Kern et al.

an organization or its IT infrastructure, incorrect policy updates and a common practice of granting permissions to freely [46] cause ACPs to deteriorate and lose quality over time [44,22]. This decay leads to inaccurate authorizations, which create security risks or prevent users from accessing resources, and a reduction in ACP maintainability, which reduces the work efficiency of policy engineers and increases their proneness to further errors [6]. Unlike their initial modeling, maintaining the quality of existing ACPs over time has received little research attention.

This work offers three contributions to address this research gap: (i) We conduct a detailed problem analysis for ACP maintenance, building on a literature survey and six expert interviews with Identity and Access Management (IAM) experts. We identify five fundamental problems relevant during ACP maintenance. (ii) We propose a framework for ACP maintenance that addresses the identified problems. It provides an Access Control Model (ACM)-independent high-level structure for maintenance activities that span from the definition of goals over the implementation of a maintenance environment to the execution of a maintenance process. (iii) We conduct a case study on ACP maintenance that instantiates the proposed framework in a real-world enterprise environment. It evaluates the proposed framework and makes ACP maintenance tangible. The remainder of this work is structured as follows: Chapter 2 introduces preliminaries and related work. Chapter 3 presents the problem analysis which characterizes the identified research gap and underlines its relevance. Chapter 4 presents the proposed framework that contributes to closing this research gap. Chapter 5 presents the case study that shows the framework's general validity. Chapter 6 discusses the results and concludes this work.

2 Background

2.1 Basic definitions and assumptions

Identity and Access Management (IAM) deals with the management of (digital) identities and the control of user access to resources. Authorizations must be defined here in order to determine which resources a user may or may not access. IAM relies on *Access Control Policies (ACPs)* [36], machine-processable rules which are automatically evaluated by an access control mechanism to make authorization decisions. The data structure of ACPs is defined by *Access Control Models (ACMs)*, with Discretionary Access Control (DAC) [38], Role-Based Access Control (RBAC) [13,37] and Attribute-Based Access Control (ABAC) [24] being among the most common. The authorizations granted by ACPs of different ACMs can be represented as an access matrix, which relates all covered subjects (users) with all covered objects (permissions) and contains the respective authorization decisions (permit or deny) as binary values. In condensed form, an access matrix can be expressed as a set of *User Permission Assignments (UPAs)*, which contains the sum of all effective permission grants defined by the ACP set as user-permission pairs. For an access control mechanism to make correct authorization decisions, ACPs must be modeled and maintained. The *IAM team*

of an organization is the group of people who are responsible for their modeling and maintaining. Depending on the organization, the IAM team can be located differently, e.g. in IT operations, risk management, IT security management, or a specialized IAM department. Besides the IAM team, there may be *policy owners* who are formally responsible for specific ACPs, e.g. because they have formal responsibility for the affected users of permissions (e.g. department heads or application administrators). In addition to owners, there are *domain experts*, i.e. people who have specific knowledge necessary for understanding and managing specific ACPs, like effects of specific permissions or required activities of employees fulfilling their work. Many established regulatory frameworks and IT security standards oblige organizations to ensure current authorizations in accordance with the principle of least privilege [33,3,27]. This may include that policy owners periodically (e.g. annually) check the correctness of existing UPAs. To do this, organizations carry out *access reviews*, a largely manual process in which responsible persons check all effective UPAs of an ACP set and try to find excessive authorizations which are then revoked [28].

2.2 Related work

Numerous publications address the initial modeling of high-quality policies. While policy engineering approaches aim to create policies from scratch in a top-down procedure [11,43], policy mining algorithms evaluate existing permission assignments to generate new policies based on them [31,47]. Policy modeling approaches of both types provide valuable assistance in the initial creation of policies. However, they do not aim to assist in maintaining or improving the quality of existing policies. Several publications propose process models or frameworks that aim to assist in ACP maintenance: Fuchs et al. propose a process model which aims to maintain high-quality roles [15]. It defines four phases in which an existing role model is assessed and updated with operations such as role shrinking, UPA cleansing, role expansion, role modeling, and hierarchy optimization. The authors have a clear organizational focus and incorporate issues such as distributed expert knowledge and maintenance priorities. However, the proposed maintenance process is limited to a "pure" RBAC. It does not guide the strategic derivation of maintenance goals or the operational involvement of domain experts. Benedetti and Mori propose a process model to include access logs into role maintenance, and a Max-SAT algorithm that evaluates them to improve role quality [7]. They specifically focus on identifying and adding missing permission assignments to the role model while keeping its complexity low. A subsequent publication extends its approach also to handle excessive permission assignments [8]. Similarly, Hummer et al. propose a process model for including access logs into policy management activities [26]. They propose to use this data to identify authorization inaccuracies in a policy set and find invalid policies automatically. Their approach does not go into the details of the subsequent maintenance activities. Instead, it suggests that policies recognized as invalid are re-mined fully automatically and recommended to a responsible human for confirmation. El Hadj et al. propose a framework that uses access logs to validate and maintain

4 S. Kern et al.

ABAC policies [19]. Their framework defines five modules that process policies and apply specific update operations to reduce complexity and remove conflicts and redundancies. Hu et al. propose a tool-based framework to support role updating [23]. The tool accepts desired UPA states as input. It generates possible role-permission and role-role relation updates that a policy administrator can apply to achieve the desired UPA state. Besides these frameworks, several ACP update algorithms were proposed [4,29]. ACP update algorithms aim to improve the quality of existing policies for a defined quality target while keeping the structure of the improved policies largely intact. They can help automate parts of ACP maintenance within a clearly defined scope but do not aim to support its technical or organizational implementation. To the best of our knowledge, no framework has been proposed to guide the maintenance of ACPs holistically in a real-world organization.

3 Problem Analysis

At the beginning of the research process we carried out a problem analysis. For this we researched common policy maintenance problems. The analysis of these problems served to better define the research gap and identify requirements for the developed framework. In the first part of the problem analysis, scientific IAM literature was examined in a structured literature survey with a scope for problems mentioned in the quality maintenance of ACPs. This grounding was then expanded with six expert interviews, in which IAM experts were asked about the procedure and known problems in ACPs maintenance. The knowledge body obtained in this way was then analyzed. Both the scientific literature and the expert interviews revealed a large number of problem aspects and examples that are difficult to survey in their entirety. We abstracted these and identified five overarching problems that have been mentioned repeatedly in literature and interviews and have a high level of validity. Table 1 shows analyzed literature that describes at least one of these problems. Table 2 shows in which expert interviews these problems were described. The remainder of this chapter describes details of the expert interviews and the five identified overarching problems.

3.1 Expert interviews

The six expert interviews were conducted according to the semi-structured interview methodology proposed by Adams [2]. We formulated a catalogue of 13 questions which were walked through with the interviewees in natural conversations. When relevant problems or details about the maintenance practice were mentioned, we deviated from this catalogue in order to pursue them more deeply. The results were transcribed and evaluated, and if anything was unclear, the interviewees were asked for clarification afterwards. In the remainder, the interviewees remain anonymous due to their employers' company policies. This enabled them to provide insight into their current challenges and issues in respect to ACPs. However, we are going to give a general classification of their

Table 1. Considered literature.

Literature	P1	P2	P3	P4	P5
L1: Jaferian et al. [28]	X	X	X	X	
L2: Puchta et al. [35]	X	X		X	X
L3: Parkinson and Khan [34]	X	X			
L4: Servos and Osborn [39]	X	X		X	X
L5: Smetters and Good [40]	X	X			
L6: Fuchs et al. [15]	X	X			
L7: Hummer et al. [26]	X	X			X
L8: Groll et al. [17]	X	X	X	X	
L9: Hill [21]		X		X	
L10: Benedetti and Mori [8]			X	X	
L11: Hu et al. [23]	X		X		
L12: Strembeck [41]		X			
L13: Xu et al. [46]	X				
L14: Xiang et al. [45]		X	X		
L15: Bauer et al. [5]	X	X			
L16: Kunz et al. [30]	X	X			X
L17: Kern et al. [29]	X	X			X

employing organizations by highlighting the approximate number of employees and managed digital identities. These numbers do not deviate from the actual numbers by more than 20%.

Expert Interview (EI)1 was conducted with an IAM governance officer of a banking group (approx. 5,000 employees and 10,000 digital identities). EI2 was conducted with an IAM governance officer and an IAM engineer of a pharmaceuticals company (approx. 15,000 employees and 30,000 digital identities). EI3 was conducted with an IAM governance officer and an IAM engineer working for an insurance company (approx. 5,000 employees and digital identities). EI4 was conducted with an IAM governance officer working for a retail company (approx. 50,000 employees and 20,000 digital identities). EI5 was conducted with the Chief Information Security Officer (CISO) of a software and consulting company (approx. 50 employees and digital identities). EI6 was conducted with two senior IAM consultants of the same company who approximated that they had completed IAM projects for a combined total of 60 customer companies. Note that some companies manage more digital identities than they have employees since they also manage access for their organizational network, like external contractors or suppliers.

All companies considered by interviews EI1-5 used RBAC as their basic authorization model. In parallel to RBAC, however, there have always been manual direct permission assignments without an intermediary role. The IAM consultants from EI6 emphasize that a pure RBAC is de facto absent in practice and is also not desirable due to the role explosion problem [14]. Moreover, all companies

6 S. Kern et al.

Table 2. Participants of the Expert Interviews (EIs).

Expert Interview (EI)	Sector	P1	P2	P3	P4	P5
EI1: IAM officer	Banking	X	X	X	X	
EI2: IAM officer, IAM engineer	Pharmaceutics	X	X	X	X	X
EI3: IAM officer, IAM engineer	Insurance	X	X	X	X	X
EI4: IAM officer	Retail	X	X	X	X	X
EI5: CISO	Software & Consulting		X	X	X	
EI6: 2 IAM consultants	Software & Consulting	X	X	X	X	X

used automation mechanisms for basic authorizations. These mechanisms permit or deny authorizations on the basis of a person’s position in the company’s organizational structure, logic-based or attribute-based assignment rules. Similarly, all companies use mechanisms to assign roles automatically to employees based on employee attributes. In addition, Segregation of Duty (SoD) rules exist with varying degrees of complexity: They range from simple 1-to-1 exclusions of two permissions over SoD matrices to very complex logic-based rule structures. Overall, the authorization structures could not be limited to a single ACM in any case. In addition, the authorization structures within a data schema were subdivided semantically: For example, roles were divided into hierarchy levels using multi-level concepts, and permissions were treated differently based on their application affiliation.

The five companies perform regular maintenance processes in the form of access reviews. In addition, reactive maintenance is carried out. The most frequently named reason are changes to the company’s organizational structure; e.g., because departments are merged or subcompanies are acquired. This typically leads to changes in the entitlement structures that are directly linked to organizational affiliation (e.g. department roles). Proactive maintenance is only carried out to a limited extent. Interviewees EI1-5 reported that isolated cases, e.g. outdated of obsolete permissions, can be conducted relatively easily, i.e. without any organizational resistance. However, they were reluctant to make changes to more complex entitlement structures, e.g. roles that could not easily be attributed to a well-defined user or permission group, due to the involved work effort and fear of errors. The IAM consultants from EI6 underlined that proactive ACP maintenance in their experience is scarce and often not carried out at all. Despite this reluctance, all interviewees emphasized that it pays off to improve entitlement structures if it can be done with a manageable amount of effort. The most frequently mentioned motivation are efficiency gains, as simpler entitlement structures allow for easier permission assignments and speed up employee onboarding, and improve entitlement maintainability. Another important motivation was maintenance decentralization, since simpler authorization structures can be better maintained by policy owners in departments without deeper

IT or IAM knowledge. Possible improvements in authorization accuracy were also often considered valuable. It also became clear, however, that compliance with regulatory requirements or supplier requirements from customers are no less important than internal motivations. Such compliance requirements in fact often represent the decisive reason for performing ACP maintenance, especially for access reviews.

3.2 Identified problems

P1. Amount and complexity of policies: The amount of ACPs in their various forms is too large to keep track of and update manually. For this reason, tool support is necessary for entitlement data overview and maintenance. This complexity was made particularly clear in the example of access reviews: Several interviewees explained that responsible policy owners often perceive this manual review of permission assignments as a "penalty work", and that it would not be enforceable without external compliance pressure. In the worst case, policy owners would blindly confirm all existing permission assignments, resulting in uncontrolled proliferation of authorizations [17]. The underlying IAM infrastructure's complexity also hampers entitlement data overview. The basic task of implementing a unified IAM data view is a nontrivial challenge because the managed permissions reside scattered in a large number of application systems. While provisioning engines and meta-data views aim to tackle this complexity, they represent only an abstraction of the underlying entitlement structures and cannot eliminate their complexity. For example, one interviewee highlighted, that their organization operates a parallel structure of in-house and cloud applications, which leads to intended redundancies in entitlement data. Specially customized meta-database views, which are supposed to provide an overview of the effective permission assignments of a user (so-called "reports"), are complex to comprehend and error-prone. Another interviewee explained that deployed data synchronization tools have malfunctioned in the past, causing errors in the entitlement data that remained unnoticed for a while. This interviewee also mentioned the problem of shadow IT, which occurs when departments set up IT applications bypassing the central IT operations: The IAM team then is not aware of the authorizations managed there and cannot maintain them [16].

P2. Distributed knowledge: The knowledge needed to manage ACPs is typically spread across an organization. IAM or IT security officers have an overview of the rough structure but find it hard assessing the effects of permissions within the applications or determining the required permissions for a specific employee. The knowledge for this typically lies with IT experts (e.g., application administrators) or domain experts (e.g., department heads). For this reason, the IAM team cannot keep authorization structures up to date on their own but rely on the cooperation with these knowledge bearers. In two interviews, experts reported that they have handed over some of their responsibility for role maintenance to IT or domain experts. Another two have stated this as a future goal. Several interviewees emphasized that it can be difficult for both, the IAM team and IT or domain experts to understand the semantic meaning

8 S. Kern et al.

behind existing permissions or ACPs. This starts with low-level problems, e.g., when permission naming is not related to any semantics (e.g. using numbers) or when descriptions documenting the business meaning are absent. The experts also highlighted an occasional absence of defined contact persons for further questions, e.g., to determine the security criticality of permissions. Great emphasis was placed on the semantic meaningfulness of authorization objects. One interviewee stated that one should be able to explain in one sentence what the content of a role or SoD rule is. Another interviewee emphasized that comprehensible entitlement structures are the central prerequisite for involving domain experts outside the IAM team in policy maintenance.

P3. Importance of business facilitation: At all interviewed companies, uninterrupted business operations are the top priority. As a result, IAM teams act very carefully not to revoke too many permissions from users, potentially causing negative business impact. When in doubt, they are often willing to put up with excess rights rather than prevent employees from doing their jobs [28,46]. For example, one interviewee reported the following typical behavior during their mover processes: When users change departments, they often execute tasks from their old department during a transition period, meaning that they might still need some permissions associated with their old department. The removal of known outdated UPAs is thus problematic and only carried out after such a transition period. The high importance of business facilitation is an obstacle to the maintenance of authorization structures and favors their proliferation.

P4. Organizational and regulatory restrictions: The interviewees unanimously reported formal hurdles in the maintenance of ACPs. These can be due to internal organizational requirements, for example, due to existing processes or "company politics" [16,25], or because of external regulations (often referred to as *regulatory compliance*). Such restrictions make it necessary to define ACP owners, for example, for all permissions within an application or for every role. Granting permissions to users or changing the structure of ACPs often requires the approval of these owners. This represents a hurdle for the maintenance of ACPs, especially if formally defined owners do not actually have the knowledge to assess a given change in a qualified manner. In addition, regulations often make more complex entitlement structures necessary. In the interviews, it was noticeable that heavily regulated financial service companies defined a larger amount and more complex SoD rules than those from less heavily regulated sectors. The IAM consultants from EI6 reported instances where additional layers were modeled into role or permission hierarchies only to accommodate responsibilities.

P5. Attribute quality: Accuracy, integrity, and timeliness of attributes of IAM-relevant data play a major role for ACP maintenance. In addition to the comprehensibility-relevant attributes of ACPs themselves, data records of users and departmental structures (e.g., HR records), as well as user accounts and permissions within individual applications, are elementary as a source of information [30]. Incorrect or outdated attributes in these data, e.g., the wrong department assignment of an employee, therefore lead to incorrect policy updates

or to a complete lack of necessary maintenance if the trust in the master data is missing. The IAM consultants from EI6 emphasized that sufficient master data quality is always a prerequisite for further data analyses and must therefore be ensured before attempting larger IAM projects (e.g., role modeling).

4 Proposed ACP Maintenance Framework

In the following we propose a framework for the maintenance of ACPs. It was developed and evaluated using the design science methodology [20] and builds on the previously presented problem analysis. The proposed framework describes activities that are necessary for the maintenance of ACPs and their successful integration in an organizational context. It defines four domains to which the maintenance activities are assigned: Governance, the IAM team, IT & domain experts and the maintenance environment. The governance domain is responsible for defining strategic goals, from which IAM maintenance activities are derived, and for reviewing the achievement of these goals. The IAM team has the operational responsibility for ACP maintenance. The IT & domain experts domain includes people with contextual knowledge that assists during ACP maintenance, as well as policy owners who must be included in maintenance activities. The maintenance environment is a collection of tools and software components that support the analysis and updating of the ACPs. Figure 1 gives a schematic overview of the four domains and the associated maintenance activities. Note that the framework in its entirety is not designed as a business process. The policy updating activities are short-term periodic tasks that are well-suited to be implemented as a process. The definition of strategic goals and quality objectives, and the implementation of analysis and updating capabilities are executed over a longer period and hence better suited for project-type organization. The remainder of this chapter presents the activities of the proposed framework.

4.1 Defining strategic IAM goals and ACP quality objectives

The governance domain defines strategic goals which serve as work basis for the IAM team. Strategic IAM goals commonly involve compliance, business facilitation, risk reduction and quality-related goals [25]. Risk reduction and business facilitation are directly related to the accuracy of ACPs, i.e. the amount of excessive and missing UPAs defined by them [6]. They are addressed by identifying which UPAs a given user *should* have, and updating the existing ACPs to correct deviations. Quality-related strategic goals, such as data quality, software quality or process quality, aim to ensure an efficient operability of IT and enable high-quality work results. The quality of ACPs significantly influences these goals: Beside accuracy, ACP quality includes maintainability, which affects administrative effort and error proneness through factors such as complexity, understandability or redundancy; as well as evaluation efficiency, which is a performance bottleneck if ACPs are evaluated in real-time [29]. Compliance goals typically overlap with the aforementioned, and may also include adherence to the principle of

sources such as human resource systems and directories), an integrated IAM data view must first be implemented. This data view bundles and normalizes the managed ACPs and associated data in an appropriate data model (e.g. [30]). The larger and more complex an IT infrastructure is, the more important it is to obtain a sufficient understanding of the managed data before the actual maintenance, e.g. through appropriate ACP visualization methods [42,12,10]. On this basis, the IAM team needs to build an overview of the total amount of ACPs managed in an organization and possible quality issues. IT and domain experts can support the IAM team by bringing in their domain knowledge when reviewing policies that affect their line of work. In return, the IAM team must enable IT and domain experts to understand the meaning of their policies and the possible consequences of changes. Once a high-level overview of the entitlement structures has been obtained, metrics can be defined to determine the ACP quality and monitor its development throughout the maintenance process.

The maintenance process can be implemented as soon as the required data has been developed. First, subsets of the ACP data must be defined whose quality is of interest and should be maintained. It is helpful to separate ACP subsets based on their data structure (e.g., different ACMs) as well as their semantic meaning (e.g., different layers in a role model, or different maintenance priorities of ABAC policies). Quality checks must then be implemented for the defined ACP subsets. A quality check comprises two elements: A check condition and a maintenance action. A check condition defines an automatically identifiable quality problem or opportunity for quality improvement. Examples of this can be a metric indicating low ACP quality, a constraint such as an SoD quality being violated, an ACP exceeding a defined timeliness (e.g. one year passed since the last review), or detectable events like the creation of a new department or IT application. Check conditions are evaluated periodically and fully automated. When need for maintenance is identified, a corresponding maintenance action is triggered.

4.3 Executing the ACP maintenance process

A quality check's maintenance action can be defined and implemented by the IAM team according to their maintenance goals. There are numerous possibilities for quality improvement and they depend on the identified quality problem and the affected ACPs: For example, a quality check that identifies excessive UPAs can trigger a permission withdrawal. The violation of an SoD rule can lead to a review of the violating ACPs, or a quality optimization algorithm can attempt to resolve identified conflicts or redundancies between multiple rules. We propose three prototypical grades of automation: Informing, recommending, and fully automated. (i) The simplest case is a purely informational request to check the identified quality problem. Such a request can be delegated to a responsible IT or domain expert who can decide based on this to manually update the affected policies or accept the quality issue. (ii) In the second level of automation, possible policy updates are generated automatically (e.g. by a quality optimization algorithm) and recommended to responsible IT or domain experts. The experts

12 S. Kern et al.

now have the option of accepting or rejecting the recommendation. It is important that the IT and domain experts understand the recommendation, i.e. its reason and its concrete effect. Tool support can enable IT and domain experts to make a qualified decision, e.g. by providing data visualization techniques or low-threshold contact options with the IAM team. If an expert decides to accept the recommendation, the update is performed and the ACP state is updated. If the decision is rejected, the rejected recommendation must be logged so that it is not proposed again. (iii) The third level is full automation. This level is limited to policies for which no human approval is required, e.g. because they are not subject to such regulations or because their security criticality is low enough. Fully automated ACP updates also require trust in the correctness of the recommendation, which must grow over time. Just like changes initiated externally, a successful maintenance action leads to an update of the ACP state and thus to an update of the measured ACP quality. The IAM team monitors the ACP quality and, based on this, adjusts existing quality checks or implements new ones.

Note that the roles of the IAM team and domain experts cannot always be clearly separated: In small organizations, the persons responsible for IAM tasks often take on the function of domain experts themselves. As the organization grows, these tasks can no longer be managed by the IAM team due to the work quantity and distributed content knowledge and must be consistently outsourced to IT and domain experts. It cannot be expected that a full ACP maintenance capability will be built immediately. The initial implementation focus should be on creating a data overview and analysis options that improve understanding of the ACPs and possible quality problems. Over time, new quality checks can be implemented sequentially in order to increase the coverage of considered quality problems and the degree of automation of the maintenance capabilities. For all updates made, complete logging of the changes to ACPs is helpful for traceability and external accountability. In addition, recording the quality development over time helps to check the success of the maintenance and offset it against the invested resources.

5 Evaluation with real-world enterprise data

To evaluate the proposed framework, it was instantiated in a case study. For this purpose, we worked with IAM practitioners of a large financial service provider, who gave us read access to the centrally managed entitlement data of their productive IT infrastructure. Our practice partners assumed the function of the governance domain, while we filled out the IAM team domain. First, we obtained an overview of the existing entitlement structures. The company uses an RBAC model with two semantic types of roles: Organization-driven *business roles* and application-specific *system roles*. In addition, there are manual permission assignments, and proprietary rules for automated permission assignment. There are also constraints that must be observed when updating the entitlement structure, including an SoD matrix that defines mutually exclusive permissions, and

application-specific restrictions for the assignment of permissions. After understanding the basic data model, we defined maintenance priorities in consultation with the practice partners. They were interested in improving the data overview and eliminating unnecessary complexity. The identification of excess authorizations and the improvement of master data quality were also of interest. At the same time, restrictions applied: First, any ACP updates had to comply with the defined constraints. Second, we had no access to domain experts or policy owners, as this would go beyond the resources provided for the case study. Third, no access logs were available, which could have provided a data basis for automated identification of excessive authorizations. Based on the available data and resources, we formulated the maintenance objective of reducing the complexity and redundancy of the ACPs. We decided to use two metrics to verify maintenance success: The Weighted Structural Complexity (WSC) defines the complexity of ACPs by summing up all contained data elements [32]. We decided to use a neutral configuration of (1,1,1,1,1). Redundancy was defined as the ratio of redundant UPAs among all UPAs [18].

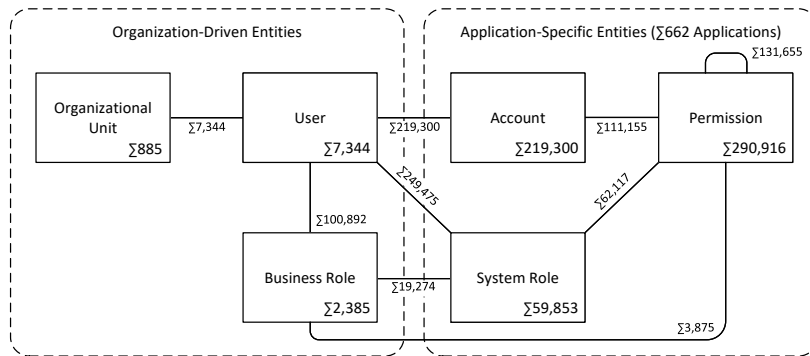


Fig. 2. Assumed data model of the case study with initial entity counts.

After agreeing on maintenance objectives and metrics we implemented the maintenance environment. Therefore we created a relational database and implemented an integrated IAM data view in accordance with [30]. It comprised the managed users, their organizational unit affiliations, user accounts and permissions, roles and the relations between these entities. Entity attributes contained rules for automatic permission assignment, SoD classes and roles and permissions, application-specific assignment constraints, and context information such as job descriptions, policy ownership definitions or security criticality flags. Figure 2 summarizes the integrated data model. In order to keep the data complexity manageable in the context of the case study, we excluded some known exceptional cases from the imported data, such as authorization assignments to external employees without a domain account, orphan accounts, or applications

14 S. Kern et al.

without centrally managed permissions. In a productive maintenance project, these exceptional cases would be considered once maintenance was established for the basic ones.

We then proceeded to implement analysis and maintenance capabilities. The data analysis capabilities included a tool-based data browsing interface, a graphic filter for the analysis of data subsets, a grid visualization for entitlement data, and quality metric calculations for selected data sets. A workflow engine enabled us to bundle data changes into change requests and delegate them to selected deciders via E-Mail. While we did not have permission to query actual domain experts, we did a proof-of-concept configuration that used ownership attributes from the imported permission data to include domain experts in the maintenance process. We then defined three subgroups within the imported ACP data subject to quality considerations: 2,385 business roles, 59,853 system roles, and 111,115 manual account-permission assignments. First quality measurements showed that the ACP set realized a total of 3,666,181 UPAs, out of which 1,134,596 (30.95%) were redundant. The initial WSC was 1,031,597. Since the business roles accounted for 62.99% of the UPAs, but only 10.01% of the complexity, we decided to minimize direct permission and system role assignments in favor of the well-maintainable business roles.

Table 3. The six quality checks implemented for the case study.

ACP Subset	Check	Condition	Maintenance action
Business roles	C1	Role without employee	Delete role
	C2	Role without permission	Delete role
	C5	All employees inherit the same system role	Assign system role to business role
	C6	All employees inherit the same permission	Assign permission to business role
System roles	C3	Redundant assignment	Revoke assignment
Manual p. ass.	C4	Redundant assignment	Revoke assignment

With the analysis capabilities in place, we proceeded to implement quality checks in two cycles. Table 3 lists all implemented checks. The first two checks were trivial: *C1* identified business roles that are assigned to no users, and *C2* identified business roles that inherit no permissions or child roles. Such "empty" roles are leftovers from past updates that bloat the ACP set and can be deleted. The checks identified 524 roles assigned to no users and 146 without permissions or child roles, with an intersection of 80 roles. Since we were surprised by the high number of results, we contacted our practice partners, who confirmed their correctness. The 590 empty business roles were hence deleted, reducing their amount to 1,795. Check *C3* identified redundant assignments of system roles to users: If a user already inherits a system role through a (well-maintainable)

business role, any direct assignment of this role was considered redundant and would be revoked. Similarly, C_4 revoked manual permission assignments if they were identified as redundant. C_3 and C_4 resulted in the deletion 24,573 direct system role assignments and 25 manual permission assignments. After the first check implementation cycle, the redundancy ratio was reduced to 25.81% and the WSC by 3.16% to 999,039.

Table 4. Quality development during the maintenance process.

	Σ UPAs	Redundant	Ratio	WSC
Initial state	3,666,181	1,134,596	30.85%	1,031,597
1 st reduction	3,412,097	880,512	25.81%	999,039
2 nd reduction	3,412,097	880,512	25.81%	932,991
Quality improv.		22.39%	5.14%	9.56%

The subsequently implemented checks C_5 and C_6 identified opportunities for structural improvement. C_5 generated recommendations to create new role hierarchy relations: If all employees of a business role inherit the same system role, the system role should be assigned to the business role as a child. By the same logic, C_6 recommended to assign permissions to a business role, unless it was already inherited by a system role with an open recommendation from C_5 . Both C_5 and C_6 omitted recommendations that would violate SoD or application-specific constraints by evaluating respective attributes of all related roles and permissions. Since C_5 and C_6 changed the structure of existing roles, we defined that the responsible role owners had to confirm these recommendations. However, since we could not contact the real role owners, we simulated this process in the workflow engine by configuring an automatic decision with an assumed acceptance probability of 80%. In the end, C_5 created 2,447 new role hierarchy relations and C_6 created 6,537 role-permission assignments, which increased the UPA coverage of the business roles. Afterwards, C_3 identified and revoked 55,282 direct user - system role assignments, and C_4 revoked 16,498 manual permission assignments, which had become redundant through these updates. At the end of the second implementation cycle, the WSC was reduced to 932,991 (-9.56% compared to the initial value) while the redundancy ratio remained at 25.81%. Table 4 summarizes the quality development. We discussed these result with the practice partners and concluded that the maintenance objectives of a substantial reduction in redundancy and complexity had been achieved. The implemented maintenance environment remains functional and can react to future changes in the underlying ACP set by triggering ACP updates with a high degree of automation. Our practice partners received the maintenance implementation and a protocol for the conducted checks and quality improvements.

6 Discussion and Conclusion

At the beginning of this work, we carried out a detailed analysis of the problem of ACP maintenance. Rigor and relevance were ensured through a structured literature search and six expert interviews. Based on this, we proposed a framework that offers guidance for the maintenance of ACPs and thus contributes to closing this research gap. The framework is not limited to a particular ACM, but provides a high-level structure for maintenance activities, spanning from the definition of quality and maintenance objectives, over the implementation of a maintenance environment, to the execution of an ACP maintenance process. We instantiated the framework in cooperation with practice partners from a large financial services company: After defining maintenance objectives and metrics, we implemented a maintenance environment and used it to significantly improve the quality of a real-world ACP data set. Due to the open structure, the proposed framework can address arbitrary quality issues with many different maintenance approaches. While the quality checks implemented for the evaluation were intentionally kept simple, they could be supplemented by more sophisticated checks to address further quality objectives or expand the maintenance for existing ones. It should be noted that the leaps in quality achieved during the evaluation can only be expected when new quality checks are carried out for the first time. Continuous execution should instead stabilize the level of quality achieved.

This work also has limitations: First, the proposed framework can only offer guidance for the identified problems *P1-4*. *P5* (insufficient attribute quality) must be addressed by data quality management measures, which are not within the scope of ACP maintenance. Due to its high level of abstraction, the framework cannot define concrete quality improvements (unlike quality optimization algorithms, for example), but serves as a template for structuring ACP maintenance in the context of an organization. During the evaluation, we could only simulate the involvement of domain experts, which limits its general validity. In addition, some constellations of the real-world ACP data set (e.g. proprietary rules for automated permission assignment) were ignored due to limited resources. Overall, we were able to show that the proposed framework has a high degree of general validity and is suitable for guiding the maintenance of ACP in a real-world environment.

Future work can address open research questions that became apparent in the course of this work. First, there are few approaches to measure or improve the human intelligibility of ACPs, which has a strong impact on their maintainability. Another open question is how excessive UPAs can be identified effectively when no access logs are available. The integration of domain experts in ACP maintenance also represents a major difficulty, for which little assistance has been provided so far. We are also not aware of any empirical data that would provide information about real ACP quality developments, for example to investigate the extent to which users accumulate excess authorizations over time. Another open challenge is the effective identification of excessive UPAs with an absence of access logs.

References

1. Owasp foundation.: Owasp top ten project. <https://owasp.org/Top10/> (2021), accessed: April 10, 2023
2. Adams, W.C.: Conducting semi-structured interviews. Handbook of practical program evaluation pp. 492–505 (2015)
3. Basel Committee on Banking Supervision: Basel accords. https://www.bis.org/basel_framework/index.htm (1988–2004), accessed: April 10, 2023
4. Batra, G., Atluri, V., Vaidya, J., Sural, S.: Incremental maintenance of abac policies. In: Proceedings of the Eleventh ACM Conference on Data and Application Security and Privacy. pp. 185–196 (2021)
5. Bauer, L., Cranor, L.F., Reeder, R.W., Reiter, M.K., Vanica, K.: Real life challenges in access-control management. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. p. 899–908. CHI '09, Association for Computing Machinery, New York, NY, USA (2009). <https://doi.org/10.1145/1518701.1518838>, <https://doi.org/10.1145/1518701.1518838>
6. Beckerle, M., Martucci, L.A.: Formal definitions for usable access control rule sets from goals to metrics. In: Proceedings of the Ninth Symposium on Usable Privacy and Security. pp. 1–11 (2013)
7. Benedetti, M., Mori, M.: Parametric rbac maintenance via max-sat. In: Proceedings of the 23rd ACM on Symposium on Access Control Models and Technologies. p. 15–25. SACMAT '18, Association for Computing Machinery, New York, NY, USA (2018). <https://doi.org/10.1145/3205977.3205987>, <https://doi.org/10.1145/3205977.3205987>
8. Benedetti, M., Mori, M.: On the use of max-SAT and PDDL in RBAC maintenance. *Cybersecurity* **2**(1) (Jul 2019). <https://doi.org/10.1186/s42400-019-0036-9>, <https://doi.org/10.1186/s42400-019-0036-9>
9. Beyond Identity: Former employees admit to using continued account access to harm previous employers (Feb 2022), <https://www.beyondidentity.com/blog/great-resignation-impact-on-company-security>
10. Colantonio, A., Di Pietro, R., Ocello, A., Verde, N.V.: Visual role mining: A picture is worth a thousand roles. *IEEE Transactions on Knowledge and Data Engineering* **24**(6), 1120–1133 (2011)
11. Das, S., Mitra, B., Atluri, V., Vaidya, J., Sural, S.: Policy engineering in rbac and abac. From Database to Cyber Security: Essays Dedicated to Sushil Jajodia on the Occasion of His 70th Birthday pp. 24–54 (2018)
12. Das, S., Sural, S., Vaidya, J., Atluri, V., Rigoll, G.: Vismap: visual mining of attribute-based access control policies. In: Information Systems Security: 15th International Conference, ICISS 2019, Hyderabad, India, December 16–20, 2019, Proceedings 15. pp. 79–98. Springer (2019)
13. Ferraiolo, D.F., Sandhu, R., Gavrila, S., Kuhn, D.R., Chandramouli, R.: Proposed nist standard for role-based access control. *ACM Trans. Inf. Syst. Secur.* **4**(3), 224–274 (aug 2001). <https://doi.org/10.1145/501978.501980>, <https://doi.org/10.1145/501978.501980>
14. Fuchs, L., Pernul, G., Sandhu, R.: Roles in information security – a survey and classification of the research area. *Computers & Security* **30**(8), 748–769 (2011). <https://doi.org/https://doi.org/10.1016/j.cose.2011.08.002>, <https://www.sciencedirect.com/science/article/pii/S016740481100099X>

18 S. Kern et al.

15. Fuchs, L., Kunz, M., Pernul, G.: Role model optimization for secure role-based identity management. In: European Conference on Information Systems (ECIS). pp. 1–15 (Juni 2014), <https://epub.uni-regensburg.de/30394/>
16. Fuchs, L., Pernul, G.: Supporting compliant and secure user handling - a structured approach for in-house identity management. In: The Second International Conference on Availability, Reliability and Security (ARES'07). pp. 374–384 (2007). <https://doi.org/10.1109/ARES.2007.145>
17. Groll, S., Kern, S., Fuchs, L., Pernul, G.: Monitoring access reviews by crowd labelling. In: Fischer-Hübner, S., Lambrinouidakis, C., Kotsis, G., Tjoa, A.M., Khalil, I. (eds.) Trust, Privacy and Security in Digital Business. pp. 3–17. Springer International Publishing, Cham (2021)
18. Guarneri, M., Arrigoni Neri, M., Magri, E., Mutti, S.: On the notion of redundancy in access control policies. In: Proceedings of the 18th ACM symposium on Access control models and technologies. pp. 161–172 (2013)
19. Hadj, M.A.E., Erradi, M., Khoumsi, A., Benkaouz, Y.: Validation and correction of large security policies: A clustering and access log based approach. In: 2018 IEEE International Conference on Big Data (Big Data). pp. 5330–5332 (2018). <https://doi.org/10.1109/BigData.2018.8622610>
20. Hevner, A., Chatterjee, S., Hevner, A., Chatterjee, S.: Design science research in information systems. Design research in information systems: theory and practice pp. 9–22 (2010)
21. Hill, L.: How automated access verification can help organizations demonstrate HIPAA compliance: A case study. J Healthc Inf Manag **20**(2), 116–122 (2006)
22. Hu, H., Ahn, G.J., Kulkarni, K.: Anomaly discovery and resolution in web access control policies. In: Proceedings of the 16th ACM symposium on Access control models and technologies. pp. 165–174 (2011)
23. Hu, J., Zhang, Y., Li, R.: Towards automatic update of access control policy. In: Proceedings of the 24th International Conference on Large Installation System Administration. p. 1–7. LISA'10, USENIX Association, USA (2010)
24. Hu, V.C., Ferraiolo, D., Kuhn, R., Schnitzer, A., Sandlin, K., Miller, R., Scarfone, K.: Guide to attribute based access control (ABAC) definition and considerations. Tech. rep., U.S. Department of Commerce (Jan 2014). <https://doi.org/10.6028/nist.sp.800-162>, <https://doi.org/10.6028/nist.sp.800-162>
25. Hummer, M., Groll, S., Kunz, M., Fuchs, L., Pernul, G.: Measuring identity and access management performance - an expert survey on possible performance indicators. In: Proceedings of the 4th International Conference on Information Systems Security and Privacy. pp. 233–240. SCITEPRESS - Science and Technology Publications (2018). <https://doi.org/10.5220/0006557702330240>, <https://doi.org/10.5220/0006557702330240>
26. Hummer, M., Kunz, M., Netter, M., Fuchs, L., Pernul, G.: Adaptive identity and access management - contextual data based policies. EURASIP Journal on Information Security **2016**(1) (Aug 2016). <https://doi.org/10.1186/s13635-016-0043-2>, <https://doi.org/10.1186/s13635-016-0043-2>
27. International Organization for Standardization: Iso/iec 27000:2013 – information technology – security techniques – information security management systems – overview and vocabulary. <https://www.iso.org/standard/54534.html> (2013), accessed: April 10, 2023
28. Jaferian, P., Rashtian, H., Beznosov, K.: To authorize or not authorize: Helping users review access policies in organizations. In: Proceedings of the Tenth USENIX

- Conference on Usable Privacy and Security. p. 301–320. SOUPS '14, USENIX Association, USA (2014)
29. Kern, S., Baumer, T., Groll, S., Fuchs, L., Pernul, G.: Optimization of access control policies. *Journal of Information Security and Applications* **70**, 103301 (2022). <https://doi.org/https://doi.org/10.1016/j.jisa.2022.103301>, <https://www.sciencedirect.com/science/article/pii/S2214212622001533>
 30. Kunz, M., Puchta, A., Groll, S., Fuchs, L., Pernul, G.: Attribute quality management for dynamic identity and access management. *Journal of Information Security and Applications* **44**, 64–79 (2019). <https://doi.org/https://doi.org/10.1016/j.jisa.2018.11.004>, <https://www.sciencedirect.com/science/article/pii/S2214212618301467>
 31. Mitra, B., Sural, S., Vaidya, J., Atluri, V.: A survey of role mining. *ACM Computing Surveys (CSUR)* **48**(4), 1–37 (2016)
 32. Molloy, I., Chen, H., Li, T., Wang, Q., Li, N., Bertino, E., Calo, S., Lobo, J.: Mining roles with semantic meanings. In: *Proceedings of the 13th ACM symposium on Access control models and technologies*. pp. 21–30 (2008)
 33. One Hundred Seventh Congress of the United States of America: Sarbanes-oxley act of 2002. <https://www.govinfo.gov/content/pkg/PLAW-107publ204/pdf/PLAW-107publ204.pdf> (2002), accessed: April 10, 2023
 34. Parkinson, S., Khan, S.: A survey on empirical security analysis of access-control systems: A real-world perspective. *ACM Comput. Surv.* **55**(6) (dec 2022). <https://doi.org/10.1145/3533703>, <https://doi.org/10.1145/3533703>
 35. Puchta, A., Böhm, F., Pernul, G.: Contributing to current challenges in identity and access management with visual analytics. In: Foley, S.N. (ed.) *Data and Applications Security and Privacy XXXIII*. pp. 221–239. Springer International Publishing, Cham (2019)
 36. Samarati, P., de Vimercati, S.C.: Access control: Policies, models, and mechanisms. In: *Foundations of Security Analysis and Design: Tutorial Lectures 1*. pp. 137–196. Springer (2001)
 37. Sandhu, R.S.: Role-based access control. portions of this chapter have been published earlier in sandhu et al. (1996), sandhu (1996), sandhu and bhamidipati (1997), sandhu et al. (1997) and sandhu and feinstein (1994). In: Zelkowitz, M.V. (ed.) *Advances in Computers, Advances in Computers*, vol. 46, pp. 237–286. Elsevier, online (1998). [https://doi.org/https://doi.org/10.1016/S0065-2458\(08\)60206-5](https://doi.org/https://doi.org/10.1016/S0065-2458(08)60206-5), <https://www.sciencedirect.com/science/article/pii/S0065245808602065>
 38. Sandhu, R.S., Samarati, P.: Access control: principle and practice. *IEEE communications magazine* **32**(9), 40–48 (1994)
 39. Servos, D., Osborn, S.L.: Current research and open problems in attribute-based access control. *ACM Comput. Surv.* **49**(4) (jan 2017). <https://doi.org/10.1145/3007204>, <https://doi.org/10.1145/3007204>
 40. Smetters, D.K., Good, N.: How users use access control. In: *Proceedings of the 5th Symposium on Usable Privacy and Security. SOUPS '09, Association for Computing Machinery, New York, NY, USA* (2009). <https://doi.org/10.1145/1572532.1572552>, <https://doi.org/10.1145/1572532.1572552>
 41. Strembeck, M.: Scenario-driven role engineering. *IEEE Security & Privacy* **8**(1), 28–35 (Jan 2010). <https://doi.org/10.1109/MSP.2010.46>
 42. Sun, W., Su, H., Xie, H.: Policy-engineering optimization with visual representation and separation-of-duty constraints in attribute-based access control. *Future Internet* **12**(10), 164 (2020)

20 S. Kern et al.

43. Verde, N.V., Vaidya, J., Atluri, V., Colantonio, A.: Role engineering: From theory to practice. In: Proceedings of the second ACM conference on Data and Application Security and Privacy. pp. 181–192 (2012)
44. Xia, H., Dawande, M., Mookerjee, V.: Role refinement in access control: Model and analysis. *INFORMS Journal on Computing* **26**(4), 866–884 (2014)
45. Xiang, C., Wu, Y., Shen, B., Shen, M., Huang, H., Xu, T., Zhou, Y., Moore, C., Jin, X., Sheng, T.: Towards continuous access control validation and forensics. In: Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security. p. 113–129. CCS '19, Association for Computing Machinery, New York, NY, USA (2019). <https://doi.org/10.1145/3319535.3363191>, <https://doi.org/10.1145/3319535.3363191>
46. Xu, T., Naing, H.M., Lu, L., Zhou, Y.: How do system administrators resolve access-denied issues in the real world? In: Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems. p. 348–361. CHI '17, Association for Computing Machinery, New York, NY, USA (2017). <https://doi.org/10.1145/3025453.3025999>, <https://doi.org/10.1145/3025453.3025999>
47. Xu, Z., Stoller, S.D.: Mining attribute-based access control policies. *IEEE Transactions on Dependable and Secure Computing* **12**(5), 533–545 (2014)

5 A Framework for Managing Separation of Duty Policies

Current status:	Published
Conference:	19th International Conference on Availability, Reliability and Security (ARES 2024), Vienna, July 30-August 02, 2024
Date of acceptance:	June 02, 2024
Full citation:	Sebastian Groll, Sascha Kern, Ludwig Fuchs, and Günther Pernul. A framework for managing separation of duty policies. In <i>Proceedings of the 19th International Conference on Availability, Reliability and Security</i> . ACM, 2024.
Authors contributions:	Sebastian Groll 50% Sascha Kern 30% Ludwig Fuchs 10% Günther Pernul 10%

Conference Description: The International Conference on Availability, Reliability and Security focuses since 2006 on rigorous and novel research in the field of dependability, computer and information security. In cooperation with the conference several workshops are held covering a huge variety of security topics.



A Framework for Managing Separation of Duty Policies

Sebastian Groll
University of Regensburg
Regensburg, Bavaria, Germany

Ludwig Fuchs
Nexis GmbH
Regensburg, Bavaria, Germany

Sascha Kern
Nexis GmbH
Regensburg, Bavaria, Germany

Günther Pernul
University of Regensburg
Regensburg, Bavaria, Germany

ABSTRACT

Separation of Duty (SoD) is a fundamental principle in information security. Especially large and highly regulated companies have to manage a huge number of SoD policies. These policies need to be maintained in an ongoing effort in order to remain accurate and compliant with regulatory requirements. In this work we develop a framework for managing SoD policies that pays particular attention to policy comprehensibility. We conducted seven semi-structured interviews with SoD practitioners from large organizations in order to understand the requirements for managing and maintaining SoD policies. Drawing from the obtained insights, we developed a framework, which includes the relevant stakeholders and tasks, as well as a policy structure that aims to simplify policy maintenance. We anchor the proposed policy structure in a generic IAM data model to ensure compatibility and flexibility with other IAM models. We then show exemplary how our approach can be enforced within Role-Based Access Control. Finally, we evaluate the proposed framework with a real-world IAM data set provided by a large finance company.

CCS CONCEPTS

• **Security and privacy** → **Access control**; *Usability in security and privacy*; *Formal security models*;

KEYWORDS

Separation of Duty, Identity and Access Management, Role-Based Access Control

ACM Reference Format:

Sebastian Groll, Sascha Kern, Ludwig Fuchs, and Günther Pernul. 2024. A Framework for Managing Separation of Duty Policies. In *The 19th International Conference on Availability, Reliability and Security (ARES 2024)*, July 30–August 02, 2024, Vienna, Austria. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3664476.3670912>

1 INTRODUCTION

The concept of Separation of Duty (SoD) is widely considered a fundamental principle in information security within organizations. SoD aims to mitigate fraud and potential conflicts of interest by



This work is licensed under a Creative Commons Attribution International 4.0 License.

ARES 2024, July 30–August 02, 2024, Vienna, Austria
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1718-5/24/07
<https://doi.org/10.1145/3664476.3670912>

dividing responsibilities and tasks among different persons. For example, a banking employee should not be able to both issue and approve a loan without at least one more person being involved in the process. Regulatory requirements such as SOX [25], HIPAA [34], BAIT [12], VAIT [6] or Basel III [4] force organisations to specify, document, and enforce SoD rules. Therefore, SoD is addressed in the academic literature concerning various contexts and research areas.

Existing research has devoted considerable attention to the specification [2, 20, 31] and enforcement [7, 9, 33] of SoD rules. However, the creation and the maintenance of SoD rules has hardly been addressed by research so far. Existing works typically assume that rules already exist or are specified by some "administrator", who knows which rules are needed. However, big organisations usually deal with thousands of employees with different functions and responsibilities leading to millions of permissions and roles in numerous departments and applications. Thus, the necessary set of SoD rules is large and complex, and the knowledge required to specify and maintain correct rules spreads throughout the organization across domain experts, managers, departmental heads, risk managers, IT specialists, etc. [18]. Once SoD rules have been initially created, they must be maintained in an ongoing effort. Similar to any set of data, SoD rules can become outdated due to changes in both external compliance requirements and the organization's internal structure (e.g. permissions or roles). Consider, for example, a new version of HIPAA, or the introduction of a new HR system in the organization. Both will lead to new requirements for the existing SoD rules. Active SoD rules must therefore be regularly reviewed, updated, and eventually de-provisioned when their validity expires.

In order to be maintainable and manageable, the meaning and purpose of each SoD rule must be clear: Where does the rule originate from? Which legal text or guideline is responsible for its existence? What happens if the legal text changes or a new one is added? Who ensures that the rule is technically implemented correctly? etc. The knowledge required to answer these questions extends throughout the entire organization. Hence, SoD rules must be created and maintained within a structured processes with clear responsibilities and knowledge holders. They need to be enriched with descriptive, human-understandable information and the knowledge spread throughout the organization must be consolidated. To the best of our knowledge, the current research does neither provide an approach for deriving and maintaining SoD rules nor does it provide an approach to enrich SoD rules with descriptive information.

This work contributes to this research gap by proposing a framework for SoD rules that are easy to maintain. Our contributions

are as follows: (i) We conduct seven expert interviews with SoD managers from large, compliance-driven companies. They serve to broaden the understanding of SoD management and to identify best practices for SoD maintenance. (ii) On this basis, we propose a framework for the creation and maintenance of SoD rules, that pays particular attention to descriptive information and comprehensibility. It defines the key stakeholders for SoD maintenance, and three types of SoD policies which are anchored in the integrated IAM data model by Kunz et al. [21]. (iii) After a conceptual definition, we formalize our easy to maintain SoD rules as an exemplary proof of concept in Role-Based Access Control (RBAC) [28] and show how they can be enforced. We then evaluate the framework by applying our approach to a real-world IAM data set provided by a large finance company. Our evaluation shows that the complexity of managing and maintaining SoD with the proposed framework is significantly lower than managing traditional SoD rules.

2 FOUNDATION

SoD rules can roughly be divided into rules with static and dynamic SoD properties [14, 31]. Static rules are generally applicable, e.g. "A user who is allowed to carry out Task A must not be allowed to carry out Task B". Dynamic SoD rules are only applicable in a specific context, e.g. at the same time, with the same object or operation. An example rule would be "A user must not carry out Task A and Task B for the same customer; however, the user can carry out Task A for one customer and Task B for another". For the sake of simplicity, we'll rely on static SoD rules throughout the remainder of the work. Whenever we refer to SoD rules, we mean static SoD rules. However, we discuss how our framework can be adapted to dynamic SoD rules in section 8.

SoD is particularly important in the context of Identity and Access Management (IAM). IAM is a domain of IT management that comprises a range of processes, technologies and policies (including SoD policies), which deal with the administration of digital identities and the provisioning of secure user access to digital resources. It commonly relies on the *principle of least privilege*, which determines that no user should have more authorizations than he or she requires to perform their duties. The structure of user authorizations is defined in IAM by Access Control Models (ACMs). While various ACMs have been developed for different scenarios, some have proven to be highly adaptive in theory and practice (especially Role-Based Access Control (RBAC) [28], Attribute-Based Access Control (ABAC) [16], and Discretionary Access Control (DAC) [27]). Regardless of the utilized ACM, the defined authorizations can be expressed as a set of User Permission Assignments (UPAs) (e.g. in the form of an access matrix [29]): A permission represents an operation and a resource (e.g. "read" and "file x"). If a permission is assigned to a user, the user is authorized to perform the access specified by it. Independent of an ACM, SoD rules can be defined as sets $sod(\{p_1, \dots, p_n\}, k)$ where p_i are permissions and k, n are positive integers such that $1 < k \leq n$ [22]. The rule states that there must not exist any set of users smaller than k that possess the permissions $\{p_1, \dots, p_n\}$ together.

While this is a very fine-grained way to specify rules, the enforcement is proven to be computationally expensive (NP-complete) [22]. SoD is extensively studied in the context of the RBAC model,

in which roles are used as intermediaries between users and permissions. A more efficient way to specify SoD rules in RBAC are Mutual Exclusive Role Policies (MER). A MER defines an amount of roles that form a toxic combination and therefore a single user is not allowed to possess together. For example the roles "Compliance Officer" and "Cashier" could form such a toxic combination, because the cashier has to carry out financial transactions and the compliance officer needs to approve them. If both roles would be assigned to the same person, this person could manipulate transactions or make false bookings. Formally, MERs are defined as sets $mer(\{r_1, \dots, r_n\}, t)$ with each r_i as a role and n, t as positive integers, such that $1 < t \leq n$ [22]. In contrast to the permission-based specification, a MER policy does not allow a user to have t or more roles from the set $\{r_1, \dots, r_n\}$.

MER policies are more coarse-grained than permission-based SoD policies but faster to enforce. It is also possible to translate permission-based SoD policies into MER policies. However, this will result in stricter rules that are less flexible. Other models that are addressed in SoD research include workflows (e.g. [10]), modelling languages (e.g. [30, 36]) or petri nets (e.g. [19]). Nevertheless, most approaches use an underlying RBAC model for SoD specification and enforcement. Note that the terms SoD policy and SoD rule are mostly used interchangeably, while the term constraint often describes SoD rules in the context of an ACM (e.g. "MER constraint"). We maintain this terminology in the course of this work.

3 RELATED WORK

Existing SoD literature covers a wide range of topics: The specification of SoD rules is studied in ACMs like RBAC [2, 20, 31] or ABAC [3, 5], in workflows [10, 23] or in petri nets [11, 40]. Other works enhance languages like UML [26, 32] or BPMN [37, 38] with SoD support. Another related research domain deals with the enforcement of SoD rules [7, 9, 33]. Despite broad coverage of the SoD concept in general, few works are dedicated to the actual creation of maintenance of SoD rules in the context of an organization. Some algorithmic approaches exist for the generation and transformation of SoD rules: Li et al. [8, 22] study the translation between permission-based SoD rules and MERs. As mentioned previously, the objective is to create MER policies that are minimally restrictive while still enforcing the underlying permission-based SoD policies. However, these works focus only on the algorithmic transformation of policies and not on the derivation of policies within organizations.

Kijsanayothin et al. [15] propose an approach to derive MER policies from a workflow. They assume an SoD violation when the same person is allowed to perform consecutive create and update actions on the same object. We argue that SoD rules for organizations need to be more fine-grained and flexible, for example also read-permissions might be relevant for SoD. Additionally, for some real-world applications it might be hard to define which permissions exactly grant update or create privileges on objects. Wolf and Gehrke [35] propose a 6-step method to derive SoD rules in an organization. Their method starts with an early analysis phase, e.g. interviewing employees in order to discover and analyze relevant processes. It continues with the specification and translation of SoD rule sets and finally describes the extraction of live data from

the applications and the enforcement of the specified SoD rules. While this work provides valuable insights into organizational aspects of SoD creation, its research scope is limited to Enterprise Resource Planning (ERP) systems. It provides neither a generic approach for deriving SoD rules nor semantics or a generic data structure for these rules. Furthermore, they do not map SoD rules to a standardized specification used in scientific literature.

The closely related concept of Access Control Policies (ACPs) describes machine-processible rules which define positive or negative authorizations to regulate which access a user is allowed to make [29]. ACPs can express SoD rules, but are not limited to them. Several works in this research realm propose to align rules with human-understandable semantics to improve their maintainability. They share the assumption that the most important factor for low management complexity of rules is that they have can be associated with a real-world concept. Such a real-world concept (e.g. "all department heads may access the following resources") can be annotated to rules via attributes which represent that concept (e.g. "function = department head").

Fuchs et al. [13] highlight the high importance of annotated business semantics for effective management or maintenance of role-based ACPs. Molloy et al. [24] propose approaches to mine role-based ACPs based either on semantically meaningful user attributes or on formal concept lattices. Similarly, Jin et al. propose to assign roles to users based on user attributes [17]. Xu defines an interpretability metric which attempts to approximate the semantic meaningfulness of role-based or attribute-based ACPs by measuring their accordance with semantically meaningful user attributes [39]. To the best of our knowledge, no scientific work exists that provides a framework for the management of SoD policies in the context of an organization, or describes a generic approach to improve their human comprehensibility.

4 SEMI-STRUCTURED EXPERT INTERVIEWS

We conducted semi-structured expert interviews to gain a better understanding of how SoD rules are modeled and maintained in large organizations. We interviewed seven experts from six different organizations in accordance with the methodology proposed by Adams [1]. We prepared a questionnaire through which we went with the interviewees in a natural conversation. If further relevant points were mentioned or ambiguities arose, we deviated from the prepared catalogue in order to investigate these in more detail. The structure of the interviews can be divided into 3 blocks, which are briefly described below:

- **B1: Introduction and General Questions**

Goal: Initiate interview. Ensure relevance of the interviewed organization.

- Short introduction.
- What is the participant's job position and and the participant's connection to SoD within the organization?
- What is the organization's size? How many employees work there and how many digital identities are managed?
- What is the main motivation for the organization to manage SoD?

- **B2: Structure and human understandability**

Goal: Discover different SoD types and their structure. Determine how meaningful information is stored.

- How is the organization's data model for SoD rules structured?
- How do they make SoD understandable to people? Do they use descriptive names, or employ other attributes?
- Which methods or tools exist in the organization to display SoDs in a human understandable way?

- **B3: Lifecycle and processes**

Goal: Clarify stakeholders and responsible parties for SoD rule creation and maintenance. Determine how meaningful information is derived.

- Who in the organization is responsible for SoD rule creation and maintenance?
- How does the organization derive SoD rules?
- How does the organization maintain SoD rules and ensure a sufficient data quality?

After conducting the interviews we structured and summarized the received answers. Naturally, there were variations in both terminology and content; for instance, some refer to roles as 'business roles' or 'organizational roles,' describing the same underlying concept. Despite these different perspectives, we could abstract several organizational structures for SoD management that are established in many or sometimes all of the considered organizations. Below we summarize significant results of the three interview blocks.

B1: The interviewed experts work on SoD projects for large organizations in different industry branches.

Organization O1 is an organization with about 17,000 employees in the engineering sector. Compared to more heavily regulated sectors, such as banking or insurance, O1 is subject to less stringent requirements. SoD rules are primarily applied to an ERP system for which a standardized set of SoD rules has been purchased. In this organization we interviewed an IAM consultant.

O2 is a large telecommunications company with about 18,000 employees. Due to both internal and external compliance requirements, SoD rules must be applied to many different application systems. They use purchased rule sets as well as self-created SoD rules. The self-created rules are derived using so-called SoD Classes, a concept we will explain further in Block B2. The interviewed expert is a risk manager.

O3 is a company with 33,000 employees. Being part of the highly regulated banking sector, it adheres to numerous legal regulations. SoD rules are created and managed using SoD Classes. Additionally, domain experts can also define pairwise mutual exclusions for roles and permissions directly. We interviewed one person working in IAM business support and one person working as a compliance manager.

O4 is an organization in the insurance sector with 7,000 employees. In the insurance sector, it is also common that compliance requirements require strong regulation through SoD. SoD rules are created and maintained solely using SoD Classes. The interviewee emphasized that the creation of SoD rules is a complex and politically sensitive task, as it relies heavily on the support of the domain experts. We interviewed a member of the IAM team.

Table 1: Overview of the seven interviewed SoD experts and their employing organizations.

Org.	Participants	Employees	Sector	Pairwise exclusive Permissions	Pairwise exclusive Roles	SoD Classes
O1	1	~17,000	Engineering	X		
O2	1	~18,000	Telecommunications	X		X
O3	2	~33,000	Finance	X	X	X
O4	1	~7,000	Insurance			X
O5	1	~300	Finance	X		X
O6	1	~4,500	Finance			X

O5 is an organization in the financial sector with about 300 employees. It is notable that even smaller companies in the banking sector are subject to strict regulations and therefore have a high administrative burden through SoD. The organization utilizes SoD Classes and pairwise mutual exclusions for permissions, that are created and maintained by the domain experts. The SoD rules have linked legal texts and descriptions explaining why the corresponding exclusion have to exist. There also exist different risk-levels for the exclusions which are useful when SoD violations are mitigated. According to O5 not all SoD violations have to be resolved: Sometimes it is possible to grant exemptions and accept the risk. The interviewed expert is an IAM consultant.

O6 is a banking company with strict regulations like the other organizations from the finance industry. SoD rules serve to avoid conflicts of interest as well as over-privileged users, especially IT administrators. They derive SoD rules using SoD Classes. The interviewed expert is a risk management officer. Table 1 summarizes the interviewed experts and their organizations. All employee numbers were rounded to ensure the anonymity of the companies.

All of the organizations use a form of RBAC with some sort of schema definition for role hierarchies. They also manage "direct" user-permission assignments which can be granted independently from the role model, e.g. through a user self service portal. Most of the organizations operate some form of automation logic which assigns roles or permissions to users based on user attributes or proprietary rules. However, for this, none of them uses formal ABAC or the eXtensible Access Control Markup Language (XACML). The amount of managed permissions ranges between 10,000 and 1,000,000. All interviewed experts named regulatory compliance as the primary driver of their SoD management.

B2: The interviewed experts described three prototypical structures for easily understandable and maintainable SoD rules: Pairwise mutually exclusive permissions, pairwise mutually exclusive roles, and the use of SoD Classes (cf. Table 1). The simplest way is the definition of pairwise mutual exclusions between two permissions or two roles. In this case, a user is not allowed to possess both permissions or both roles at once. In contrast to the MERs defined in literature (cf. Chapter 2), which allow the definition of larger sets and a minimum number of allowed roles, the mutual exclusions in the interviewed organizations were always limited to two roles or two permissions. Such pairwise mutual exclusions are easier to understand and manage. The interviewees also stated that it is important that every mutual exclusion requires a proper and up-to-date name and description annotation to remain human-understandable. Despite their simplicity, the management of mutual

exclusions has some shortcomings: They quickly amount to large data sets, causing unnecessary complexity.

Another way the experts mentioned to define SoD rules is the use of so-called SoD Classes. SoD Classes represent groupings of tasks, functions or responsibilities in an organization that are designed to conflict with each other. The SoD Classes have a name with a semantic meaning and may also have a description for further information. They are assigned to roles and permissions and a user is not allowed to possess roles or permissions with conflicting SoD Classes. According to the interviewed experts, the main benefit of SoD Classes lies in their simple management and human understandability: Typically, large organizations manage hundreds of thousands of roles and permissions, while managing only 15 to 50 SoD Classes. Another advantage is that SoD Classes are well-suited for addressing the challenge of knowledge distribution among various stakeholders regarding SoD: Legal experts can define the SoD Classes and conflicts, while technical or Domain Experts can determine the matching SoD Class for a permission or role. This becomes evident in the next section.

B3: The experts named various stakeholders which are responsible for the SoD rule creation and maintenance: Normally, the responsibility for overall SoD management lies within an IAM team. However, Domain Experts, like owners of application systems, permissions, or roles, may also be responsible, especially in large systems. For example, the IAM Team may be responsible for the entirety of SoD rules and SoD Classes, but creating mutual exclusive permissions or roles or assigning SoD Classes to permissions may be the responsibility of the permission and roles owners. Also there usually is some kind of IAM governance department, that is responsible for risk assessment and the interpretation of legal texts and compliance requirements. The creation of SoD rules typically follows one of three prototypical methods: (i) *Buying complete rule sets*, (ii) *Defining SoD classes* and (iii) *individual fine-grained rule definition*.

(i) Buying a complete rule set is a common approach to ensure compliance with regulatory requirements. Such rule sets can for example be obtained from auditing companies and cover the permissions ranged in a specific kind of target system. Since an organization does not acquire the know-how necessary to maintain these rules when purchasing, a dependency on the supplier can arise: Without regular follow-up purchases, the purchased rule sets can out-date and lose their validity over time. From a purely compliance-driven point of view, however, they are the "safe bet": First, buying standardized rules sets that are confirmed to be compliant (at the time of purchase) gives an organization some degree

of security that they will meet regulatory requirements for the covered application systems. Second, such standardized rule sets do not require to be human understandable since they are not maintained by the organization itself, but merely enforced. Therefore, those rule sets usually do not contain lots of human-understandable information. The rule sets are typically limited to one application and do not support organization-wide SoD policies. Additionally, they can also be expensive and are not available for every application, e.g. applications that are not used widely or in-house developments.

(ii) The use of SoD Classes involves multiple stakeholders: The IAM governance that defines the SoD Classes and various Domain Experts that are owners of permissions or applications. The SoD Classes are derived by the IAM governance by evaluating compliance regulations and using industry-standards. They do not usually change much over time and hence require little maintenance. The assignment of the SoD Classes to the permissions is carried out by a permission or application owners, depending on the respective application and how the application is managed. In order to keep the rules up to date, regular reviews can be carried out.

(iii) Finally, all organizations use some sort of fine-grained SoD management. This can include pairwise mutual exclusive roles and permissions, but also proprietary formats, e.g. logic-based rules or custom evaluation scripts. Managing these fine-grained SoD rules requires a deep understanding of the underlying authorization structures and has to be done by role, permission or application owners or administrators. These rules also need to be reviewed periodically to maintain their accuracy and timeliness.

5 DERIVING A FRAMEWORK FOR MANAGING AND MAINTAINING SOD

Based on the insights we gathered from the conducted expert interviews, as well as on existing SoD and IAM literature, we derived a framework for managing and maintaining SoD policies. The framework consists of three components: (i) The relevant stakeholders and their responsibilities, (ii) the three proposed SoD rule types along with the respective processes for rule creation and maintenance and (iii) the integration of the proposed SoD rule types into existing ACMS. The developed framework is presented in the remainder of this chapter. Figure 1 depicts the described tasks and responsibilities of the different stakeholders in our framework.

5.1 Stakeholders and responsibilities

In order for SoD policies to be well-understandable, they need to hold a real-world meaning. Such semantic meaningfulness does not arise from the structure of policies alone, but must be incorporated into them during their creation and maintenance. Not everyone involved in the management of SoD rules has all the necessary information available. Therefore, we first define three types of stakeholders that were described during the interviews:

- **IAM Governance / Risk Management** (*Called IAM Governance in the following*): Responsible for complying with laws and regulatory requirements. Ensures that audit requirements are met. Manages the IAM at a high level, specifies processes and carries out risk assessments. Knows the laws and legal requirements, but does not necessarily have

a deeper understanding of individual roles, permissions and IT applications.

- **IAM Team** Conducts IAM projects, is responsible for the implementation and enforcement of SoD policies and rules. Understands permissions and roles on a conceptual and technical level but cannot always assess their business implications. The IAM team size can vary from few persons up to a whole IAM Department.
- **Domain Experts** Are responsible for one or more IT application systems, processes or organizational units. Typical positions are department heads, business owners or technical owners like system admins. Domain experts know the tasks of users and the effects of permissions and roles in their responsibility in great detail. However, they do not have an overarching view of the IAM structure and may not necessarily be aware of possible interactions with other domains.

Since the necessary knowledge for creating and maintaining SoD rules is distributed across these domains, collaboration between the involved stakeholders is necessary. The IAM team serves as an intermediary between the IAM governance and the domain experts.

5.2 The SoD Matrix: A tool to reduce complexity

The SoD Matrix is a central control tool of our framework for SoD management. It acts as an intermediary between the IAM team, IAM governance and the domain experts. Preliminary for its creation is the definition of SoD classes, which are represented in the SoD matrix as rows and columns. A SoD class is an entity that describes a real-world concept relevant to SoD. An example for SoD classes could be "Payment Transactions" or "Internal Audit".

The SoD Matrix defines a pairwise exclusion whenever two SoD classes are incompatible with one another. Due to its Matrix form, the SoD Matrix is intuitively understandable and suitable for visualization. Figure 2 displays a (reduced) SoD Matrix as managed by a real-world finance company. The definition and maintenance of the SoD Matrix (and SoD classes) is carried out by the IAM governance through reading and interpreting legal texts and security standards. The IAM Governance has the required expertise and knowledge for this task. A deeper understanding of the applications or their permissions and roles is not required. For example, a member of IAM Governance could read in a compliance standard that functions related to payment transactions must not be combined with functions dealing with compliance and approval. The IAM governance employee would create the SoD classes "Payment Transactions" and "Internal Audit" as well as an exclusions between the two SoD classes in the SoD matrix.

In order to enforce the SoD matrix on the application level, the SoD classes need to be assigned to the permissions of the respective applications. This task is carried out by the domain experts, as it requires in-depth knowledge of the permissions, particularly regarding which actions can be performed with which permissions. At this point, it is crucial that the SoD classes have meaningful names and descriptions to be understood by the domain experts. For example a permission that allows a user to edit a financial transaction on an ERP system would be assigned to a SoD class "Payment Transactions" by a domain expert. Note that many permissions are

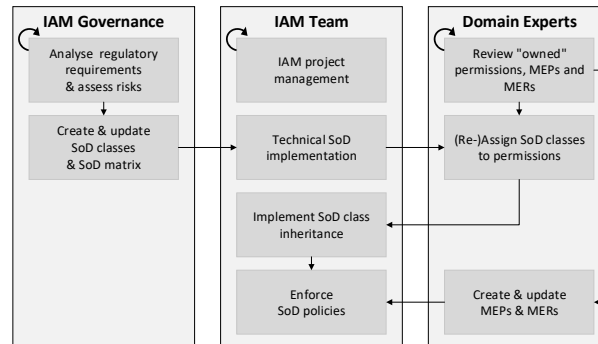


Figure 1: Responsibilities and activities of the proposed framework.

not relevant for SoD (e.g. "WiFi access"). All permissions without a matching SoD class are assigned a neutral SoD class which can never cause a SoD conflict with another SoD class.

Applying the SoD Matrix to all permissions results in a list of pairwise mutual permission exclusions. The IAM team must bring the SoD Matrix and the permissions with SoD classes together and enforce the resulting exclusions. These can be implemented according to any established SoD specification (cmp. section 2). Maintaining the resulting (technical) SoD rules no longer requires involving governance or domain experts. In return, governance and domain experts can keep the SoD components in their responsibility up-to-date on their own, without having to rely on the knowledge of others. In addition to the IAM project organization required for the creation of a SoD Matrix, the IAM team must provide the technical infrastructure and organize regular reviews. The described responsibilities and tasks can be seen in Figure 1.

5.3 Fine-grained exclusions with pairwise MERs and MEPs

	Market	Market Follow-Up	Audit	Risk Controlling	Accounting	Legal	Compliance	Trade	Payment Traffic	Fund Mgt.
Market	Grey	Red	Red	Red	Green	Green	Red	Green	Green	Green
Market Follow-Up	Red	Grey	Red	Red	Green	Green	Red	Green	Green	Green
Audit	Red	Red	Grey	Red	Green	Green	Red	Green	Green	Green
Risk Controlling	Red	Red	Red	Grey	Green	Green	Red	Green	Green	Green
Accounting	Green	Green	Green	Green	Grey	Green	Red	Green	Green	Green
Legal	Red	Red	Red	Red	Green	Grey	Red	Green	Green	Green
Compliance	Red	Red	Red	Red	Red	Red	Grey	Red	Green	Green
Trade	Green	Green	Green	Green	Green	Green	Green	Grey	Green	Green
Payment Traffic	Red	Red	Red	Red	Green	Green	Red	Red	Grey	Green
Fund Mgt.	Green	Green	Green	Green	Green	Red	Red	Green	Green	Grey

Figure 2: Matrix visualization of SoD classes and their pairwise mutual exclusions ("SoD Matrix").

Note: Red cells mark a pairwise exclusion, while green cells indicate that the two SoD classes are not conflicting. The depicted SoD Matrix defines 31 pairwise SoD class exclusions.

The SoD matrix enables well-maintainable and compliant, but coarse-grained SoD rules. In real-world applications, more precise SoD rules may be necessary. For example, two specific permissions (or roles) may form a toxic combination, but be included in compatible SoD classes. This can happen for technical reasons, especially since permissions function slightly differently in each application. Also there may not be a suitable SoD class for some roles.

The knowledge about these exclusions lies with the domain experts. With pairwise Mutually Exclusive Roles (MER) and pairwise Mutually Exclusive Permissions (MEP) they can create direct exclusions, in addition to the SoD matrix. In contrast to MERs specified in scientific literature, we propose using MERs and MEPs that require a mandatory description for each exclusion explaining its respective reason. Furthermore, they have a fixed cardinality of 2, which makes pairwise MERs an instance of the default RBAC MER: $mer_2\{r_1, r_2\} \equiv mer\{\{r_1, r_2\}, 2\}$. The processes for creating and maintaining MERs and MEPs are performed by the domain experts and organized by the IAM team. Again the IAM team is responsible for providing a suitable infrastructure and organizing regular reviews.

5.4 Data model integration

Existing IAM literature defines numerous ACMs designed for different scenarios and requirements. Among the most commonly used ACMs are RBAC, ABAC and the somewhat older but still relevant DAC model. Alongside these theoretical data structures exist industry standards for identity data exchange, SSO and authorization delegation. Since a precise specification of the proposed SoD policies in a significant amount of ACMs and standards exceeds the scope of this work, we chose to anchor it in the conceptual IAM data model proposed by Kunz et al [21]. It provides a generic view over the central entities defined in the ACMs, e.g. RBAC and ABAC, and in the standards LDAP, SAML, SPML, OAuth, SCIM and XACML. To the best of our knowledge, it is the most generic IAM data model proposed in scientific literature. A more detailed formalization of the proposed SoD policies for the RBAC model is provided in section 6.

The conceptual data model defines a digital identity as a representation of a human user. An account is a representation of a user

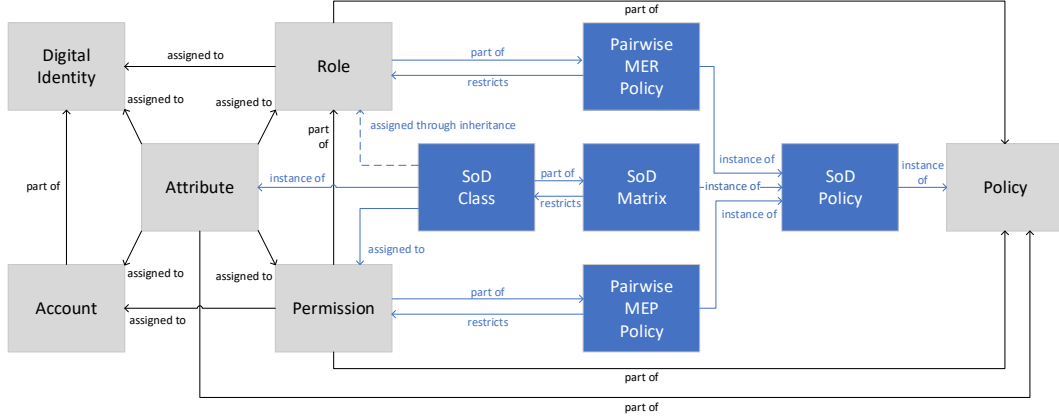


Figure 3: Anchoring of proposed SoD policies in the conceptual IAM data model by Kunz et al. [21].

Note: Grey entities were adopted from the original conceptual IAM model, while blue entities are introduced in the proposed SoD framework.

within an application system and can hence be part of a digital identity. Permissions enable users to execute certain actions, commonly expressed as a combination of a resource to be accessed and an access operation (e.g. modify files in a certain directory). Roles are semantic entities that bundle a set of permissions. They are valid organization-wide and can hence be assigned directly to a digital identity. Policies represent sets of rules that define authorizations, such as ABAC, SPML or XACML. Attributes represent meta-data that is assigned to any of the described entities. An example of a user attribute would be their job title or the department they work in. The original data model also defines a context entity, which we omit because it is not relevant for the scope of this work.

We extend the conceptual model by Kunz et al. [21] to cover the proposed SoD policies. SoD policy is a specific type of policy that constrains the authorization structure by defining mutually exclusive entitlements. A violation of this constraint represents a conflict that must be resolved. We define three kinds of SoD policies: Pairwise MERs, MEPs and the SoD matrix (along with its SoD classes). Pairwise MER and MEP policies, as specified before, define simple pairs of mutually exclusive roles or permissions. They are hence linked to exactly two permissions (MEP) or roles (MER) each and restrict their assignment to a digital identity: Any pair of mutually exclusive permissions or roles that is inherited by a single digital identity constitutes a violation of the defining policy which must be resolved to restore compliance.

The SoD matrix defines mutually exclusive SoD classes. The SoD classes are assigned to permissions as an attribute. Every permission has exactly one SoD class assigned, which may also be the neutral SoD class that indicates no SoD relevance. Roles inherit SoD classes transitively from their permissions. We constrain the model by demanding that roles are *homogeneous*, i.e. must not inherit permissions with more than one SoD class other than the neutral one. Each pair of mutually exclusive SoD classes in the SoD matrix indicates that no digital identity may inherit permissions from both of these classes. Figure 3 summarizes the SoD policies from the proposed framework, anchored in the conceptual IAM model.

6 POLICY FORMALIZATION AND ENFORCEABILITY IN RBAC

We extended the conceptual IAM model by Kunz et al. to anchor the SoD policies in a generic data model. In the course of this chapter, we will provide a more detailed formalization in the RBAC model to demonstrate its feasibility. We then show that the proposed SoD policies can be translated into classic MER constraints. We chose RBAC because it is a well-known ACM in SoD literature as well as in the organizations of the interviewed SoD experts. A practical evaluation is provided subsequently in section 7.

6.1 Model formalization

We adopt the preliminaries and notations from the RBAC SoD formalizations by Li et al. [22] for sakes of consistency.

Preliminaries. Let U be the set of all users, R the set of all roles and P the set of all permissions. An RBAC state γ is a 3-tuple $\langle UA, PA, RH \rangle$, where $UA \subset U \times R$ is the set of all user-to-role assignments, $PA \subset R \times P$ is the set of all role-to-permission assignments, and $RH \subset R \times R$ is the set of all role-to-role assignments. RH^* is the reflexive, transitive closure of RH and a partial order among roles in R . The functions $auth_roles_\gamma[u]$ and $auth_perms_\gamma[u]$ determine the roles and permissions that a user effectively inherits and are defined as:

$$auth_roles_\gamma[u] = \{r \in R \mid \exists r_1 \in R [(u, r_1) \in UA \wedge (r_1, r) \in RH^*]\}$$

$$auth_perms_\gamma[u] = \{p \in P \mid \exists r_1, r_2 \in R [(u, r_1) \in UA \wedge (r_1, r_2) \in RH^* \wedge (r_2, p) \in PA]\}$$

Definition 1. A pairwise MER policy mer_2 is a specific type of MER constraint with fixed value of $n = t = 2$. It defines a pair of mutually exclusive roles

$$mer_2\{r_1, r_2\}$$

with $r_1, r_2 \in R$ and $r_1 \neq r_2$. A pairwise MER policy is satisfied by an RBAC state γ if no single user inherits both roles defined by it

as mutually exclusive:

$$sat_{mer_2}[\gamma] \triangleq \forall u \in U(|auth_roles_\gamma[u] \cap \{r_1, r_2\}| < 2)$$

Definition 2. A pairwise MEP policy mep_2 defines a pair of mutually exclusive permissions

$$mep_2\{p_1, p_2\}$$

with $p_1, p_2 \in P$ and $p_1 \neq p_2$. It is satisfied by an RBAC state γ if no single user inherits both permissions defined by it as mutually exclusive:

$$sat_{mep_2}[\gamma] \triangleq \forall u \in U(|auth_perms_\gamma[u] \cap \{p_1, p_2\}| < 2)$$

Definition 3. An organization manages a finite set of $I + 1$ SoD classes with \emptyset being the *neutral SoD class*:

$$SOD_CLASSES = \{\emptyset, class_1, \dots, class_I\}$$

Every permission p is assigned a SoD class $assigned_class_p[p] \in SOD_CLASSES$. By default, any permission has $assigned_class_p[p] = \emptyset$, unless it was assigned a non-neutral $class \in \{class_1, \dots, class_I\}$. The set of all permissions which share a SoD class $class$ is determined by the function

$$perms[class] = \{p \in P | assigned_class_p[p] = class\}$$

A role inherits SoD classes transitively from its child roles and permissions:

$$trans_class[r] = \{class \in SOD_CLASSES | \exists r_1 \in R((r, r_1) \in RH^* \wedge (r_1, p) \in PA \wedge assigned_class_p[p] = class)\}$$

The effectively assigned SoD classes are consequently:

$$assi_class_r[r] = \begin{cases} \{\emptyset\} & \text{if } trans_class[r] = \{\emptyset\} \\ & \text{or } trans_class[r] = \{\emptyset\} \\ trans_class[r] \setminus \{\emptyset\} & \text{otherwise} \end{cases}$$

We call an RBAC state γ *homogeneous* if no role inherits more than one non-neutral SoD class:

$$homogeneous[\gamma] \triangleq |assi_class_r[r]| = 1 \forall r \in R$$

The set of all roles which share a SoD class $class$ is determined by the function

$$roles[class] = \{r \in R | class \in assi_class_r[r]\}$$

Definition 4. A user inherits all SoD classes of her assigned roles, except the neutral one:

$$assigned_class_u[u] = \{class \in SOD_CLASSES \setminus \{\emptyset\} | \exists r_1 \in R[r_1 \in auth_roles_\gamma[u] \wedge assi_class_r[r_1] = class]\}$$

Definition 5. A SoD matrix is a set of pairwise mutual SoD class exclusions. A pairwise mutual SoD class exclusion mec_2 defines a pair of mutually exclusive SoD classes

$$mec_2\{class_1, class_2\}$$

with $class_1, class_2 \in SOD_CLASSES \setminus \emptyset$ and $class_1 \neq class_2$. It is satisfied by a homogeneous RBAC state γ if no single user inherits roles which are assigned both SoD classes defined by it as mutually exclusive:

$$sat_{mec_2}[\gamma] \triangleq \forall u \in U(|assigned_class_u[u]| < 2)$$

Note that we intentionally limit the use of SoD policies to homogeneous RBAC states to improve role semantics. While the model

formalization would be equally applicable for non-homogeneous RBAC states, we argue that an RBAC state can be made homogeneous relatively easy, which improves the overall policy understandability and maintainability (cmp. sections 5 and 8).

6.2 Policy Enforcement

In the following we show that the proposed SoD policies are compatible with established SoD specifications for RBAC. In RBAC, SoD rules are commonly specified as MER constraints $mer(\{r_1, \dots, r_n\}, t)$, with each r_i as a role and n, t as positive integers, such that $1 < t \leq n$. If t or more roles are assigned to a user, the rule is violated. Formally, an RBAC state satisfying an MER constraint is denoted as follows [22]:

$$sat_{mer}[\gamma] \triangleq \forall u \in U(|auth_roles_\gamma[u] \cap \{r_1, \dots, r_n\}| < t)$$

Thus, to show how the proposed SoD policies can be enforced in RBAC, we have to show how the three types of SoD policies (pairwise MEP, MER and mutually exclusive SoD classes) can be translated to MER constraints, as defined above. This is trivial for the pairwise MER policy since it can be defined as MER with $n = t = 2$:

$$mer_2\{r_1, r_2\} \equiv mer(\{r_1, r_2\}, 2)$$

A pairwise MEP policy is equivalent to the set of pairwise MER policies that defines pairwise exclusions for all roles which inherit the mutually exclusive permissions transitively:

$$mep_2\{p_1, p_2\} \equiv \{mer_2\{r_1, r_2\} | \forall r_1 \in R[\exists r \in R((r_1, r) \in RH^* \wedge (r, p_1) \in PA), r_2 \in R[\exists r \in R((r_2, r) \in RH^* \wedge (r, p_2) \in PA)]]\}$$

A pairwise mutual SoD class exclusion is equivalent to the set of pairwise MER policies that defines pairwise exclusions for all roles with the specified SoD classes:

$$mec_2\{class_1, class_2\} \equiv \{mer_2\{r_1, r_2\} | \forall r_1 \in roles[class_1], r_2 \in roles[class_2]\}$$

Both pairwise MEPs and mutual SoD class exclusions can be translated into a set of pairwise MER policies. Since any pairwise MER policy can be expressed as a classic MER constraint, any SoD matrix can be translated into established MER constraints and enforced as such.

7 EVALUATION

We cooperated with a large financial services provider to evaluate the proposed framework. An IAM expert from the company also took part in the expert interviews (see organization O6 in table 1). The conceptual overview of stakeholders and responsibilities was generally agreed upon as it was synthesized from the expert interviews. We hence defined two evaluation criteria: (i) *Technical applicability*: Can we confirm that the proposed SoD matrix and mutual pairwise exclusions are compatible with existing SoD specifications and can be enforced as such? (ii) *Simplified policy management*: Can we confirm that the framework provides a complexity reduction compared to established MER policies?

We were allowed to analyze anonymized, productively used SoD entitlement data of the organization. O6 employs around 4,500 people and manages around 10,000 digital identities. Similar to the

other companies in the expert interviews, the banking group organizes SoD policies through SoD classes and a SoD matrix. The permissions in the data set are organized in the form of an RBAC model with a flat role hierarchy: The roles are called "business roles". Business roles are valid in the context of the entire company and bundle users based on a functional assignment, e.g. "customer support" or "liquidity forecast". The permissions are organized in two layers: "system entitlements" and "technical permissions". Technical permissions are fine-grained permissions that are only assigned to system entitlements. SoD classes are only assigned to system entitlements, which makes the technical permissions irrelevant for our use case. Hierarchy relations between two business roles or two system entitlements are not allowed. We analyzed a total of 14 SoD classes and about 8,000 system entitlements.

We implemented two algorithms that allow us to generate MERs based on the SoD matrix, SoD classes and system entitlements. While we could not include them in this work due to space constraints, we provide them on a GitHub repository along with sample data to execute them¹. The first algorithm collects SoD classes from permissions and transitively writes them to the inheriting roles. It guarantees the roles' SoD class homogeneity and marks roles that would otherwise inherit more than one SoD class. If a role would inherit more than one SoD class we assign \times , which represents an invalid SoD class state. These violations can be addressed by splitting the permissions of the role into multiple roles. The second algorithm transforms a SoD Matrix and a set of homogeneous roles into equivalent MER constraints.

By applying the algorithms on the data, we assigned at least one SoD class to 209 roles. Out of these, 5 had more than one SoD class. Based on these assignments, we generated a total of 12,295 MERs. Table 2 summarizes the results. We validated the results with the IAM experts from O6. The experts confirmed the correctness of our results and the usefulness of our approach. They pointed out that the effort and error rate in managing SoD classes is significantly lower compared to managing MERs. MERs, however, can be enforced more easily and are immediately compatible with common IAM systems. On this basis we conclude that the evaluation criteria were met:

(i) *Technical applicability*: We were able to transform a SoD matrix and permissions assigned with SoD classes of a real-world banking group into classic MERs. This means that a SoD matrix can be enforced as such. The enforcement of pairwise MERs and MEPs is trivial and the correctness of our translation was confirmed by the IAM experts of O6.

(ii) *Simplified policy management*: The quantifiable data complexity required to manage a SoD matrix is significantly lower than the data complexity for equivalent MERs. The real-world data set contained 14 SoD classes, 32 pairwise SoD class exclusions, 274 permissions with a non-neutral SoD class and 209 roles with at least one non-neutral SoD-class, amounting to a total of 529 entities that need to be managed. In contrast, an equivalent set of pairwise MERs amounts to 12,295 entities. In theory, the MERs could be algorithmically optimized to cover the same exclusions with fewer MERs that have a higher cardinality. However, these MERs would be optimized for low data volume only and would not

¹https://github.com/sgroll/semantic_sod

Table 2: Output of algorithm execution on real world data

Basic data	
Roles	2,494
Permissions	7,972
Role-permission assignments	18,692
SoD class data	
SoD classes	14
Pairwise SoD class exclusions (SoD Matrix)	32
Permissions with non-neutral SoD class	274
Roles with non-neutral SoD class	209
Algorithm results	
Resulting MERs	12,295
Homogeneity violations	5

represent a human-understandable real-world concept. Note also that individual stakeholders have to manage fewer than 529 entities: Governance employees only need to maintain the SoD matrix with 14 SoD classes and 32 pairwise exclusions, and domain experts only manage permissions, of which only 274 have a non-neutral SoD class.

8 DISCUSSION & CONCLUSION

This work proposed a framework for the creation and maintenance of semantically meaningful SoD policies. The work is grounded on seven expert interviews with SoD practitioners from the industry. In these interviews we attempted to gain an understanding and summarize best practices for SoD management in large and compliance-driven organizations. The developed framework aims to provide structures for semantically meaningful SoD which remain aligned with established practices.

The framework starts with a definition of SoD stakeholders and their responsibilities. It then defines three types of SoD policies which are designed to be well-maintainable and easily understandable. Finally, the defined SoD policies are anchored in an integrated IAM data model to generalize their scope. After presenting the proposed framework, this work provides a formalization of the three defined SoD policy types in the well-established RBAC model. We show that the policies can be translated into classic MER constraints and evaluate this with a real-world data set, thus proving that they can be evaluated and enforced efficiently.

The proposed SoD policies are easily maintainable and understandable because of three characteristics: (i) All managed policies have a semantic meaning derived from a real-world requirement. (ii) The SoD matrix decouples the creation of SoD exclusions from the management of permissions, which eliminates the need for collaboration between stakeholders from different knowledge domains. (iii) The quantifiable data complexity of SoD exclusions managed via the SoD matrix is significantly lower than the complexity of equivalent MERs. Our work currently only applies to rules with static SoD properties. This can be addressed by developing dynamic SoD classes: Similar to dynamic SoD rules, a user may possess permissions of conflicting SoD classes. When permissions of one of the two classes are used within a specified context (e.g. time, operation, customer etc.), the permissions of the other class cannot be used

ARES 2024, July 30–August 02, 2024, Vienna, Austria

Groll et al.

in this context. We leave a detailed specification and analyses of dynamic SoD classes for future work.

ACKNOWLEDGMENTS

The research leading to these results was supported by the German Federal Ministry of Education and Research as part of the DEVISE project (<https://devise.ur.de/>).

REFERENCES

- [1] William C. Adams. 2015. Conducting semi-structured interviews. *Handbook of practical program evaluation* (2015), 492–505.
- [2] Gail-Joon Ahn and Ravi Sandhu. 2000. Role-Based Authorization Constraints Specification. 3, 4 (nov 2000), 207–226. <https://doi.org/10.1145/382912.382913>
- [3] Nuray Baltaci and James Joshi. 2020. A Constraint and Risk-aware Approach to Attribute-based Access Control for Cyber-Physical Systems. *Computers & Security* 96 (05 2020), 101802. <https://doi.org/10.1016/j.cose.2020.101802>
- [4] Basel Committee on Banking Supervision. 2011. *Basel III - A Global Regulatory Framework for More Resilient Banks and Banking Systems*. Bank for International Settlements.
- [5] Khalid Bijon, R. Krishnan, and R. Sandhu. 2013. Constraints specification in attribute based access control. *SCIENCE 2* (01 2013), 131–144.
- [6] Bundesanstalt für Finanzdienstleistungsaufsicht. 2024. *Versicherungsaufsichtliche Anforderungen an die IT (VAIT)*.
- [7] David W Chadwick, Wensheng Xu, Sassa Otenko, Romain Laborde, and Bassem Nasser. 2007. Multi-session Separation of Duties (MSoD) for RBAC. In *2007 IEEE 23rd International Conference on Data Engineering Workshop*. 744–753. <https://doi.org/10.1109/ICDEW.2007.4401062>
- [8] Hong Chen and Ninghui Li. 2006. Constraint Generation for Separation of Duty (SACMAT '06). Association for Computing Machinery, New York, NY, USA, 130–138. <https://doi.org/10.1145/1133058.1133077>
- [9] Weihe Chen, Zhu Tang, and Shiguang Ju. 2008. Enforcement of Spatial Separation of Duty Constraint. In *2008 The 9th International Conference for Young Computer Scientists*. 2108–2114. <https://doi.org/10.1109/ICYCS.2008.223>
- [10] Jason Crampton. 2005. A Reference Monitor for Workflow Systems with Constrained Task Execution (SACMAT '05). Association for Computing Machinery, New York, NY, USA, 38–47. <https://doi.org/10.1145/1063979.1063986>
- [11] Marcelo Antonio de Carvalho and Paulo Bandiera-Paiva. 2017. Evaluating ISO 14441 privacy requirements on role based access control (RBAC) restrict mode via Colored Petri Nets (CPN) modeling. In *2017 International Carnahan Conference on Security Technology (CCST)*. 1–8. <https://doi.org/10.1109/CCST.2017.8167833>
- [12] Deutsche Bundesbank. 2017. Bankaufsichtliche Anforderungen an die IT (BAIT). Available at <https://www.bundesbank.de/de/aufgaben/bankenaufsicht/einzelaspekte/risikomanagement/bait/bankaufsichtliche-anforderungen-andie-it-598580>.
- [13] Ludwig Fuchs, Michael Kunz, and Günther Pernul. 2014. Role model optimization for secure role-based identity management. (2014).
- [14] Virgil D Gligor, Serban I Gavrilă, and David Ferraiolo. 1998. On the formal definition of separation-of-duty policies and their composition. In *Proceedings. 1998 IEEE Symposium on Security and Privacy (Cat. No. 98CB36186)*. IEEE, 172–183.
- [15] Rattikorn Hewett, Phongphun Kijsanayothin, and Aashay Thipse. 2008. Security analysis of role-based separation of duty with workflows. In *2008 Third International Conference on Availability, Reliability and Security*. IEEE, 765–770.
- [16] Vincent Hu, David Ferraiolo, D. Kuhn, A. Schnitzer, Knox Sandlin, R. Miller, and Karen Scarfone. 2014. Guide to attribute based access control (ABAC) definition and considerations. *National Institute of Standards and Technology Special Publication* (01 2014), 162–800.
- [17] Xin Jin, Ram Krishnan, and Ravi Sandhu. 2012. A role-based administration model for attributes. In *Proceedings of the first international workshop on secure and resilient architectures and systems*. 7–12.
- [18] Sascha Kern, Thomas Baumer, Ludwig Fuchs, and Günther Pernul. 2023. Maintain High-Quality Access Control Policies: An Academic and Practice-Driven Approach. In *IFIP Annual Conference on Data and Applications Security and Privacy*. Springer, 223–242.
- [19] Konstantin Knorr and Harald Weidner. 2001. Analyzing Separation of Duties in Petri Net Workflows. In *Information Assurance in Computer Networks*, Vladimir I. Gorodetski, Victor A. Skormin, and Leonard J. Popyack (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 102–114.
- [20] D. Richard Kuhn. 1997. Mutual Exclusion of Roles as a Means of Implementing Separation of Duty in Role-Based Access Control Systems. In *Proceedings of the Second ACM Workshop on Role-Based Access Control (Fairfax, Virginia, USA) (RBAC '97)*. Association for Computing Machinery, New York, NY, USA, 23–30. <https://doi.org/10.1145/266741.266749>
- [21] Michael Kunz, Alexander Puchta, Sebastian Groll, Ludwig Fuchs, and Günther Pernul. 2019. Attribute quality management for dynamic identity and access management. *Journal of Information Security and Applications* 44 (2019), 64–79. <https://doi.org/10.1016/j.jisa.2018.11.004>
- [22] Ninghui Li, Mahesh V. Tripunitara, and Ziad Bizri. 2007. On mutually exclusive roles and separation-of-duty. *ACM Trans. Inf. Syst. Secur.* 10, 2 (may 2007), 5–es. <https://doi.org/10.1145/1237500.1237501>
- [23] Jan Mendling, Karsten Ploesser, and Mark Strembeck. 2008. Specifying Separation of Duty Constraints in BPEL4People Processes. In *Business Information Systems*, Witold Abramowicz and Dieter Fensel (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 273–284.
- [24] Ian Molloy, Hong Chen, Tiancheng Li, Qihua Wang, Ninghui Li, Elisa Bertino, Seraphin Calo, and Jorge Lobo. 2008. Mining roles with semantic meanings. In *Proceedings of the 13th ACM symposium on Access control models and technologies*. 21–30.
- [25] One Hundred Seventh Congress of the United States of America. 2002. Sarbanes-Oxley Act of 2002. Pub. L. No. 107-204, 116 Stat. 745.
- [26] Indrakshi Ray, Na Li, Robert France, and Dae-Kyoo Kim. 2004. Using Uml to Visualize Role-Based Access Control Constraints (SACMAT '04). Association for Computing Machinery, New York, NY, USA, 115–124. <https://doi.org/10.1145/990036.990054>
- [27] Pierangela Samarati and Sabrina Capitani de Vimercati. 2000. Access control: Policies, models, and mechanisms. In *International school on foundations of security analysis and design*. Springer, 137–196.
- [28] Ravi S Sandhu. 1998. Role-based access control. In *Advances in computers*. Vol. 46. Elsevier, 237–286.
- [29] Ravi S Sandhu and Pierangela Samarati. 1994. Access control: principle and practice. *IEEE communications magazine* 32, 9 (1994), 40–48.
- [30] M. E. Shin and G. Ahn. 2001. Role-Based Authorization Constraints Specification Using Object Constraint Language. In *2012 IEEE 21st International Workshop on Enabling Technologies: Infrastructure for Collaborative Enterprises*. IEEE Computer Society, Los Alamitos, CA, USA, 157. <https://doi.org/10.1109/ENABL.2001.953406>
- [31] R.T. Simon and M.E. Zurko. 1997. Separation of duty in role-based environments. In *Proceedings 10th Computer Security Foundations Workshop*. 183–194. <https://doi.org/10.1109/CSFW.1997.596811>
- [32] Karsten Sohr, Gail-Joon Ahn, Martin Gogolla, and Lars Migge. 2005. Specification and Validation of Authorisation Constraints Using UML and OCL. In *Computer Security – ESORICS 2005*, Sabrina de Capitani di Vimercati, Paul Syverson, and Dieter Gollmann (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 64–79.
- [33] Karsten Sohr, Tanveer Mustafa, Xinyu Bao, and Gail-Joon Ahn. 2008. Enforcing Role-Based Access Control Policies in Web Services with UML and OCL. In *2008 Annual Computer Security Applications Conference (ACSAC)*. 257–266. <https://doi.org/10.1109/ACSAC.2008.35>
- [34] U.S. Department of Health & Human Services. n.d.. American Health Insurance Portability and Accountability Act. <https://www.hhs.gov/hipaa/index.html>. Accessed: 2023-06-10.
- [35] Patrick Wolf and Nick Gehrke. 2009. Continuous compliance monitoring in ERP systems-A method for identifying segregation of duties conflicts. *Wirtschaftsinformatik Proceedings 2009 39* (2009).
- [36] Christian Wolter and Andreas Schaad. 2007. Modeling of Task-Based Authorization Constraints in BPMN. In *Business Process Management*, Gustavo Alonso, Peter Dadam, and Michael Rosemann (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 64–79.
- [37] Christian Wolter and Andreas Schaad. 2007. Modeling of Task-Based Authorization Constraints in BPMN. In *Business Process Management*, Gustavo Alonso, Peter Dadam, and Michael Rosemann (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 64–79.
- [38] Christian Wolter, Andreas Schaad, and Christoph Meinel. 2007. Deriving XACML Policies from Business Process Models. In *Web Information Systems Engineering – WISE 2007 Workshops*, Mathias Weske, Mohand-Said Hacid, and Claude Godart (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 142–153.
- [39] Zhongyuan Xu. 2014. *Mining meaningful role-based and attribute-based access control policies*. Ph. D. Dissertation. State University of New York at Stony Brook.
- [40] Benyuan Yang and Hesuan Hu. 2023. Analysis of Authorization Constraints via Integer Linear Programming. *IEEE Transactions on Knowledge and Data Engineering* 35, 3 (2023), 2258–2271. <https://doi.org/10.1109/TKDE.2021.3124271>

6 Monitoring Access Reviews by Crowd Labelling

Current status:	Published
Conference:	18th International Conference on Trust, Privacy and Security in Digital Business (TrustBus 2021), Virtual Event, September 27–30, 2021
Date of acceptance:	June 09, 2021
Full citation:	Sebastian Groll, Sascha Kern, Ludwig Fuchs, and Günther Pernul. Monitoring access reviews by crowd labelling. In Simone Fischer-Hübner, Costas Lambrinoudakis, Gabriele Kotsis, A. Min Tjoa, and Ismail Khalil, editors, <i>Trust, Privacy and Security in Digital Business</i> , pages 3–17, Cham, 2021. Springer International Publishing.
Authors contributions:	Sebastian Groll 50% Sascha Kern 30% Ludwig Fuchs 10% Günther Pernul 10%

Conference Description: In the domain of digital businesses, concerns are raised regarding the lack of trust in electronic procedures and the extent to which information security and user privacy can be ensured. In answer to these concerns, the 17th International Conference on Trust, Privacy and Security in Digital Business (TrustBus'20) will provide an international forum for researchers and practitioners to exchange information regarding advancements in the state of the art and practice of trust and privacy in digital business.



Monitoring Access Reviews by Crowd Labelling

Sebastian Groll^{1,2}(), Sascha Kern^{1,2}, Ludwig Fuchs², and Günther Pernul¹

¹ University of Regensburg, Universitätsstraße 31, 93053 Regensburg, Germany
sebastian.groll@wiwi.uni-regensburg.de

² Nexis GmbH, Franz-Mayer-Straße 1, 93053 Regensburg, Germany

Abstract. Access reviews, i.e. the periodical security audit of access privileges, are a basic compliance and IT-security requirement for medium- and large-scale organizations. Assessing the quality of the reviewer's decisions ex-post can help to analyse the effectiveness of the measure and to identify structural or organizational shortcomings. Yet, current studies merely focus on improving the decision-making process itself. This paper develops a method for assessing the decision quality of access reviews by applying a solution from the crowd sourcing research realm. In order to achieve this, the problem of assessing decision quality of access reviews is generalized. It is shown that the abstract problem can be mapped to the problem of assessing the quality of crowd tagging decisions. Subsequently, an applicable solution of this research area is applied to access reviews. Furthermore, the selected approach is optimized to meet the specific challenges of access review data.

Keywords: Access reviews · Decision quality · Crowd sourcing · Identity and Access Management · Compliance

1 Introduction

Access reviews are a crucial task for Identity- and Access Management (IAM) and a basic compliance and IT security requirement for medium- and large-scale organizations [17, 18]. While security concerns like insider threats represent an intrinsic motivation, external regulations and IT security standards¹ are an important driver for access reviews in practice. As access privileges are a subject to constant change, access reviews must be performed on a regular basis. They are typically executed by reviewers that hold responsibility for either an access privilege or an entity that it is assigned to. During an access review, a responsible human being manually inspects assignments of access privileges and decides if they are legitimate. If they are not, the privilege assignments are thereupon revoked. Note that access reviews are not used to evaluate the access

¹ Relevant examples are the Sarbanes-Oxley Act [2], Basel III [3], the European General Data Privacy Regulation [1], ISO 27002 [4] and the BSI Grundschutz [5].

4 S. Groll et al.

to resources in real-time but rather to inspect and certify an existing structure of access privilege assignments. The effective and efficient execution of access review campaigns remains a major challenge that is leveraged by technical and organizational complexity, time pressure and the sheer amount of review subjects [20]. In fact, large organizations may have numerous user accounts assigned to hundreds of thousands of access privileges spreading over dozens of application systems [13]. As a result, the manual inspection of all assignments, and the careful assessment of their legitimacy, are a time-consuming and error-prone procedure.

In the course of this work, we propose a generalized problem formulation for the quality definition of an access review decision. We then draw a link to the field of crowd sourcing and map the defined problem instance to the research area of crowd labelling. Subsequently, we adapt a well-known approach from existing literature for evaluating the quality of crowd labelling decisions and apply it to the evaluation of access review decisions. Consequently, we present a generalized data mining algorithm that can identify access review decisions with low quality.

In order to ensure both, relevance, and rigor, we follow the principles of the Action Design Research (ADR) methodology and continuously evaluate and improve our approach working together with practitioners on a real data set. The practitioners are responsible for managing the access review campaign rather performing the access reviews themselves. We cooperate with two companies, one IT security company with an expertise in IAM data analytics, and a large European company that performs company-wide access reviews on a regular basis. This work provides four major contributions: (1) We derive an abstract definition of an access review using a conceptual IAM model, (2) we show that research and methodologies of the crowd sourcing realm can successfully be applied to access review decision making and (3) adapt and improve a crowd labelling approach to access review quality assessment, thus contributing to both, the IT security and the crowd sourcing research realm. Finally, following the principles of ADR, (4) we evaluate and improve our approach by conducting a real-life project evaluation.

2 Theoretical Background

In the following we outline the most important definitions for our work in the IAM research area. Subsequently we give an overview of the relevant literature regarding IAM, access reviews and crowd labelling.

According to Pfitzman and Hansen [31] the term identity is a set or subset of attributes or characteristics of an entity that makes this entity uniquely identifiable among other entities. Relying on this definition, we define the term employee as referring to a uniquely identifiable real-world person working for a specific company. Following Pfitzman and Hansen [31], we use the definition of digital identity to derive the term account: An (user) account is unique in a specific application system and represents its employee in this context.

In order to manage the privilege assignments, a wide range of so-called access control models have been published. Sandhu et al. [33] propose the role-based

access control model (RBAC) inducing the role as intermediary between an employee and his privileges. Using roles can significantly reduce the number of assignments and therefore increase comprehensibility [14]. Fuchs and Preis [15] argue that there are multiple semantic types of roles like business roles, organizational roles or IT-roles². Shortcomings of RBAC like the lack of fine-grained access control [6, 36] or continuously increasing role sizes [26] and numbers [25] led to the development and improvement of attribute-based access control (ABAC) concepts [20, 21].

The aim of access reviews is to inspect and certify the privilege assignments resulting from an implemented access control model. This might be done by reviewing certain permission assignments in an application or by reviewing conceptual assignments e.g. business roles which do not really exist on the application. Jaferian et al. [20] identify key challenges of access reviews and deal with the question of how to design an access review application in order to help the reviewers to make meaningful decisions. They design and implement a tool that is supposed to aid reviewers by enhancing the user interface. However, their focus lies on designing and evaluating the UI rather than evaluating the decision quality of the reviewers. Hill [17] presents a case study that showcases the introduction of software-supported access reviews and other access control management measures in a healthcare enterprise. Bobba et al. [7] discuss the design of tools that are supposed to aid domain administrators in the execution of access reviews. To the best of our knowledge, there are no scientific publications that explicitly cover the challenge of measuring access review decision quality in a structured manner.

Crowd sourcing is a collaborative and distributed problem-solving activity, where many workers join to solve various kinds of tasks. Examples for crowd sourcing projects are Galaxy Zoo³, Amazon's Mechanical Turk⁴, where users can publish tasks for crowdsourcing, or even collective projects like Wikipedia [24]. In Galaxy Zoo, for example, people work together in order to manually classify or tag different types of galaxies, like elliptical or spiral galaxies. Crowd sourcing comprises all sorts of tasks a community of crowd workers can perform.

In the following we focus on crowd labelling, which is a part of crowd sourcing. Crowd labelling is typically used to create labels for a large amount of data, e.g. to create input for supervised machine learning algorithms [16]. For example, above-mentioned Galaxy Zoo is a crowd labelling project. While crowd labelling has become a multi-million-dollar industry, the evaluation and improvement of crowd-sourced labels remains a crucial challenge (Lease 2011). The quality of labelling decisions is influenced by the labelers' expertise [27, 30, 38], their motivation [19, 32], organizational circumstances and other factors [11, 35]. As a result, crowd labelling may suffer from poor labelling quality or "spammers" [23] and thus numerous approaches to rate the quality of labels or crowd workers

² In the following we will refer to the term business role in order to emphasize the application- and technical- independent character of roles.

³ <https://www.zooniverse.org/projects/zookeeper/galaxy-zoo/>.

⁴ <https://www.mturk.com/>.

6 S. Groll et al.

have been published [8, 16, 22, 30]. Other approaches aim to improve the quality of labels in a generic way, e.g. by sorting out low quality labellers [10], or programmatically generating “gold labels” (i.e. correct labels that can be used to test the accuracy of labellers) [30].

2.1 Research Method

In order to design an ensemble artifact that addresses and satisfies organizational needs and ensures the rigor required for scientific contributions, we utilize the ADR method [34]. ADR relies on a close cooperation between researchers and practitioners and provides a framework for the dynamic interaction between the two parties. It defines four main stages containing principles to follow during their respective execution.

In the (1) Problem Formulation Phase the research opportunity is identified. The specific problem can be triggered by researchers, practitioners, or end-users. One crucial and challenging task is to define “the problem as an instance of a class of problems” [34]. This ensures that a solution for a more generalized problem is created and may extend the range of applicable theories. In this study the definition of a generalized problem enabled us to apply research approaches from the crowd sourcing realm on a problem of the IAM research area. Another principle of this phase is to design an artifact by utilizing theories. Finally, this phase also contains setting up an ADR team, which consists of practitioners and researchers, securing the long-term commitment of the practitioners as well as defining the roles and responsibilities of the team members. In the (2) Building, Intervention and Evaluation (BIE) phase the artifact is shaped in an iterative process. The practitioners and end-users implicitly take part in the design process by continuously evaluating the artifact. Moreover, the members of the ADR team benefit from their diverse expert knowledge. For example, our team consists of IAM researchers, IAM consultants and practitioners, respectively having different perspectives and insights on IAM and access reviews. While executing phase 1 and phase 2 the researcher simultaneously executes the (3) Reflection and Learning Phase. Thereby the researcher continuously transforms the gained knowledge to a broader class of problems. Therefore, our solution does not only improve the quality assessment of access reviews but is applicable for the broader field of crowd labelling as well. In the final phase (4) Formalization of Learning the researcher outlines the results achieved during the development of the artifact and its application to the organizational context. Due to the context-specific nature of ADR projects, a generalization of the concrete solution is necessary.

3 Problem Formulation

3.1 Practice-Inspired Research

For this study we worked together with two companies: The first one is a small IT security company, which offers tools and consulting services for mid-sized and

large organizations in the field of IAM. The provided services include the tool-based preparation and execution of access reviews. The second company is a large European organization with more than 15.000 employees. This company is forced to conduct company-wide access reviews on a regular basis due to existing compliance regulations. Our research was triggered by questions raised by customers of the partnering IAM security company like “How can I detect abnormalities in the access review process?”, “How can I find out if the reviewers are doing their job carefully enough?” or “How can the correctness of decisions be ensured?”. Thus, summarizing the requests and discussions we had with customers of the IT security company, we formulate the initial research question as follows: “*How can low-quality review decisions in an access review be identified ex-post?*”

Defining the roles of the practitioners [9], we decided that their responsibilities contain the provision of (anonymized) realworld IAM data (i.e. the access review data), the provision of contextual knowledge, as well as the regular evaluation of the developed artifact during the execution of the BIE cycle.

3.2 Defining Access Review Decision Quality as an Instance of a Class of Problems

According to Sein et al. [34], a “critical element of the problem formulation phase is defining the problem as an instance of a class of problems” [34]. This step is necessary to draw and contribute foundations and methodologies from and to the knowledge base. Thus, a funded and more general conceptual definition of the term “decision” in the context of access reviews. To achieve this, we adapt the existing conceptual IAM model proposed by Kunz et al. [26]. Their generic IAM model is a derivation of several relevant IAM standards (including RBAC and ABAC) and technologies containing and relating, amongst other things, employees, which are called digital identities in their work, accounts, permissions and business roles (Fig. 1). Please note that Fig. 1 represents a simplified version of the original IAM model, specifically adapted to the requirements of representing access reviews and review decisions⁵. Employees refer to the real persons working for a company. They are represented by their accounts on the specific application systems of the organization. The accounts themselves are assigned to permissions representing fine-grained authorization objects to access services offered by these systems. Permissions may be structured hierarchically, i.e. owning a parent permission leads to the possession of additional child permissions. In contrast to permissions, business roles are less technical entities within IAM environments which are typically used to represent organizational or functional roles of employees like “Accountant” or “Sales Representative” [15]. They are used to bundle permissions from different application systems to grant employees access to all application systems needed for the daily work or for executing specific tasks or business processes. In contrast

⁵ Therefore, we removed “Context” and “Policy” from the model because these entity types do not influence access reviews. We also renamed the term “(Digital) Identity” to “Employee” in order to fit our definitions and naming conventions, however, the meaning remains the same.

8 S. Groll et al.

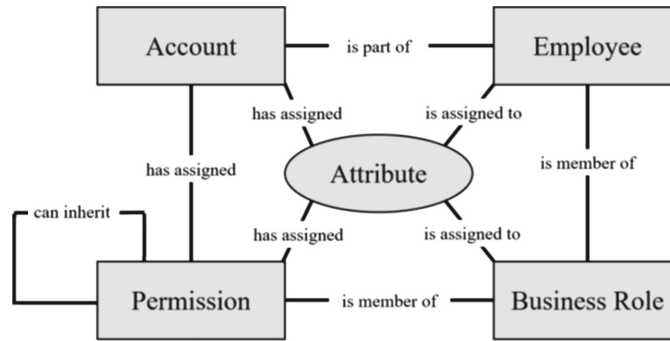


Fig. 1. Conceptual IAM model derived from Kunz et al. [26]

to accounts, business roles are directly assigned to employees, underlining their application-independent focus. Each of the entities described by Kunz et al. [26] may possess a set of attributes, which contains metadata of the respective entity. A typical attribute for an employee may be a job title or the location she is currently working at. Utilizing the conceptual IAM model of Kunz et al. [26] as well as our experiences from practical industry projects, we differentiate access reviews in two different main classes: (1) An entity itself, respectively its attributes, may be reviewed and (2) the relations between the entities, i.e. the edges in Fig. 1, may be reviewed. In the remainder of our practical analysis, we focus on the second class of access reviews (relationship reviews), however we argue that our research in general can be easily applied on the first class as well. Further generalizing, a decision d in an access review contains two entities (e.g. an account and a permission) and the decision whether these two entities should remain assigned to each other or not. Every entity e comprises a set of attributes x_e . Combined with the expert and contextual knowledge of the reviewer, these attributes form the foundation of the decision to be made. An access review decision by a reviewer r concerning the relation between two entities can further be expressed as $d_r(e1, e2) = \{x_{e1}, x_{e2}, y\}$ where y may be true or false, assuming the labels “keep relation” or “remove relation”. Thus, an access review decision made by reviewer r comprises a set of features from the two entities and the information if these entities should remain linked or not. Finally, an access review with n decisions can be represented as a set $AR = \{d_i | i \in 1, \dots, n\}$. Because we are especially interested in excessive assignments our research problem can be described as assessing the quality of every d_i with the label $y = \text{“keep relation”}$ to find suspicious and possibly wrong decisions.

4 Theory-Ingrained Artifact

According to the principles of ADR, we draw from the foundation of existing literature, to create the initial design of an artifact informed by theories. Therefore, we show that our generalized problem can be transferred to the research area of crowd sourcing and we apply selected approaches from this area. We also use well-established visualization techniques from IAM research to visualize our results.

The quality assessment problem of access reviews derived in the previous chapter can be mapped to the problem researchers are facing when assessing the

labelling quality of crowd workers. In an access review campaign, every decision $d(e1, e2)$ is a labelling task, with the labels “keep relation” and “remove relation”. Like our previously defined vector of entity attributes x_e , the data which must be labelled in crowd labelling comprises several features and is typically expressed as a feature vector. And like one of the main problems of crowd labelling, our aim is to assess the quality of the decisions (i.e. labels).

Visualization is a key component of data analytics that can leverage the insights gained from otherwise less comprehensible data sets and make them accessible to non-technical experts. We suggest extending the access grid that was introduced by Meier, Fuchs and Pernul [28] for this sake as it is a well-proven interactive visualization technique that is tailored for displaying IAM data in an easily understandable manner. The reviewed entities are displayed as rows, and the referenced entities are displayed as columns. Each cell hence depicts the relationship between two entities: If an assignment exists between these two entities, the cell is filled; otherwise it is left blank. An assignment that exists, but was not subject of the review, is filled grey (i.e. assignments that deliberately were not part of the access review campaign). The color of a cell indicates a reviewer’s decision: If an assignment was approved i.e. a positive decision was made, the cell is green; a negative decision is marked by a orange filling. If a positive decision is detected by our algorithm’s outcome and therefore potentially has low quality, then a violet border is placed around the otherwise green cell (See Fig. 2).

4.1 Designing the Initial Artifact

To design the initial artifact, we utilize theories of the crowd sourcing domain and access grid visualizations to assess the quality of a set of given access review decisions. When mapping access reviews to crowd sourcing, there are several challenges, causing access reviews to be a more specific problem in the crowd sourcing domain: (1) Access reviews, as defined in the previous section, only consist of binary decisions, while crowdlabelling may have more than one possible label e.g. more types of galaxies to classify. (2) During access reviews, one decision is usually made by one reviewer (e.g. a responsible department owner of an employee or a responsible business role owner) whereas in typical crowd labelling tasks, several reviewers label the same data. Therefore, approaches like majority vote, which are often used in crowd labelling [8, 22, 37], cannot be applied easily. (3) Following the principle of least privilege [12, 33], the quality of positive decisions is more important in access reviews than the quality of negative decisions.

4.2 Applying Crowd Sourcing Theories to Access Reviews

Based on these assumptions, we selected the machine learning approach of Geva and Saar-Tsechansky [16] for our initial artifact design. It aims to create a ranking of decision makers based on their decision quality and is capable of handling binary decisions. One of its key assumptions is that every decision is created by exactly one reviewer. The underlying research question is: “how can we rank

10 S. Groll et al.

workers by the relative quality of their decisions without resorting to the acquisition of additional and potentially costly, peer-review by other experts?” [16]. The key idea is to generate a so-called Pseudo Ground Truth (PGT) for every decision, i.e. an estimation of the correct y . This PGT can be compared with the decisions a decision maker has made to estimate her decision quality and to create a ranking of the decision makers. In order to create the PGT, a so-called base model is generated for every decision maker. The base model is trained with decision data, i.e. all decisions $d_r = \{x, y\}$ where r is the respective decision maker. These calculated base models can predict a decision for a given feature vector x , even if the decision maker, on whom the model is based, never made this decision. In other words, the base models are able to simulate reviewers’ decisions. Following this concept, it is possible to create the PGT for a decision of a certain decision maker by conducting a vote with the base models of all the other decision makers. Note that the base models are necessary because of our assumption that every decision was only carried out by one decision maker and therefore it is not possible to use the decision data of the other decision makers directly. If a decision would have been carried out by multiple decision makers, a simulation using base models would not be necessary, because we could use the concrete decisions themselves. After being generated, the PGT is used to create a decision quality score for every decision maker by dividing the decisions, where the PGT and the decision maker agree, by the total amount of decisions that the decision maker has made. Geva and Saar-Tsechansky [16] claim that the precision of the decision quality score can further be improved by introducing the concept of confidence within the PGT and applying it as weight to every decision. To calculate the confidence value, the result of the vote from the different base models is used. Finally, the calculated decision quality scores of the decision makers can be used to create a ranking of decision makers regarding their decision quality. This general approach can be adopted to our problem: Throughout the process of an access review, the reviewer represents the decision maker according to the model of Geva and Saar-Tsechansky [16]. The feature vector x is represented by the attributes of the two entities that the reviewed assignment refers to, e.g. an account and a certain permission in an application system. The actual decision (“keep relation” or “remove relation”) is the label y . Usually, several attributes for different entity types (e.g. employee, permission, etc.) are available in an identity management system and it is a critical and complex task to select a meaningful set of attributes. For example, the employee’s gender usually has no influence on her permissions but the function she holds within the organization might heavily influence the assignment of certain permissions.

5 Building, Intervention and Evaluation

5.1 Mutually Influenced Roles

During the design and learning process, there was ongoing knowledge and feedback communication between the ADR team and the partnering IT security company. Firstly, the IT security company provided the data analysis platform

that enabled the ADR team to explore the customer data and to understand the review decisions in detail and as-a-whole. Secondly, the ADR team required contextual knowledge to understand individual review decisions in the real data set, and to determine ex-post whether a decision was erroneous or not. This context knowledge was acquired through workshops with responsible IT staff of both partnering companies. Information provided by the second industry partner regarding their organizational structures, attribute semantics, as well as responsibilities of reviewers proved to be a valuable addition to the evaluation process.

5.2 Access Review Campaign Data

The industry partner executing the access review provided the ADR team with a real data set of a company-wide, quarterly access review campaign. The access review campaign was exclusively executed for account to permission assignments. The assignments were presented to the responsible reviewers in the form of employee to permission assignments to simplify their user interface and foster user acceptance. The delivered data set⁶ comprised a total of 71.464 decisions regarding 10.891 employees and 1.623 permissions and was used both for the development of the artifact and for its evaluation. The provided decisions included 3.461 decisions that had no label (i.e. reviewer decision) asserted. This is in fact not unusual for an access review campaign of this scale due to organizational circumstances (e.g. illness of reviewers, technical challenges, etc.). Unlabelled decisions were omitted by our algorithm and did not affect the results. Table 1 shows the key figures of the data set. The ADR team used to work on various subsets of the data with a meaningful filter selection, i.e. decisions regarding employees of a single department and permissions of a single application system. This way, it was possible to use existing contextual knowledge, or to pose specific inquiries to representatives of the company to gather their expertise and to understand the origin of the observed decisions.

Table 1. Key data of the evaluation data sets.

Key data	Data set
Employees to review	10.891
Reviewers	855
Application systems	13
Permissions	1.623
Positive/Negative decisions	67.116 (93,9%)/887 (1,2%)
Not decided (reviewer did not make a decision)	3.461 (4,8%)
Total decisions to make	71.464

⁶ The data set was gathered by using exports of the company's human resource system as well as exports of the different application systems.

12 S. Groll et al.

5.3 Reciprocal Shaping

After implementing the initial artifact, the ADR team carried out the approach on the real data set in a first BIE cycle. However, the generated PGT comprised almost only positive decisions. Further analyzing the results highlighted three major reasons: Firstly, there was a very strong bias towards positive decisions in the data set with 93.9% of the provided decisions being positive, which was expected because most permission assignments are correct. The second discovered limitation is that a lot of the calculated base models had to make decisions that they were not qualified for. For example, a base model generated for a reviewer who had only made decisions regarding employees in the marketing department and permissions for application system “A” now was queried about a permission from application system “B” being assigned to an employee of the software development department. Such unqualified base models, which do not share any attributes with the currently processed decision, tend to guess a positive decision. The reason is, that most of the decisions used to generate the model are positive themselves. The third issue was noise that was generated by base models that were based on very few or even only one single decision. Their decision making was inflexible and used to overrule correct negative decisions just as unqualified base models did. As a result of these issues, the base models almost always voted for positive decisions, and a completely positive PGT was generated. To address these challenges, we extended the voting process of the base models in a way that the votes are no longer considered equally meaningful. The underlying assumption is that a base model is more successful in identifying errors correctly than another if it is based on more similar decisions.

We represent this qualification of a base model $b(r)$, based on reviewer r , for a decision d by calculating an experience score $e_{d,b}$. Consequently, a decision made by this base model during the voting process is weighted with this experience. Let B be the set of all base models and B_d^+ be a subset of B with all base models that generate a positive label for decision d . Then the quality of a decision d^r made by reviewer r is the sum of the experience scores of B_d^+ , divided by the sum of the experience scores of B . The base model based on reviewer r itself is excluded from this score. The quality can then be expressed as:

$$q_{d,r} = \frac{\sum_{b^+ \in B_d^+ \setminus \{b(r)\}} (e_{d,b^+})}{\sum_{b \in B \setminus \{b(r)\}} (e_{d,b})}$$

We then define a quality threshold t that represents a minimum quality score which needs to be undercut to generate a finding for a suspected false decision. This way it is possible to adjust the output sensitivity of the algorithm. Formally, the decision d made by reviewer r is suspicious in terms of its quality, when $q_{(d,r)} < t, 0 \leq t \leq 1$. The quality threshold t can be adjusted to increase or decrease the amount of generated negative decisions and hence determines the detection sensitivity of the algorithm. While we argue that there is no determined optimum value as the amount of generated negative decisions depends on a data set’s positive decision bias, we achieved good results with a threshold

of approximately 0.6⁷. As base model decisions are made based on the assigned entities' attributes x_{e1} and x_{e2} , we define that a base model has a higher experience than another if it has made more decisions on entities with equal attributes. We hence calculate a matching attribute score $\text{attr}_{d,b}$ for each base model decision by summing up all attributes of all decisions of the current base model, i.e. the decisions of the reviewer the base model is based on, that are equal to the current decision's attributes x_{e1} , x_{e2} . This matching attribute score is the basis of our introduced experience score. In order to normalize the sum of matching attributes, we used a sigmoid function, which is often used in artificial neural network research for similar normalization problems [29]. The sigmoid function ranks base models with few decisions very low, but still allows all base models with a certain "maturity" to be taken into consideration in the weighted vote. The experience score is hence defined as $e_{d,b} = \text{sigmoid}(\text{attr}_{d,b})$. With respect to the last ADR phase (Reflection and Learning) we emphasize that this improvement can be easily adopted to the generalized crowd labelling approach by simply enhancing the voting process of the base models with the method described in this chapter. Our approach can be helpful when dealing with data with an extreme bias. However, we argue that it might also be useful when dealing with non-biased data sets, as including the experience score decreases the influence of unqualified base models and increases the influence of qualified base models.

5.4 Evaluation

The final evaluation data set was selected by filtering the review campaign data to decisions regarding permissions of one large application system. It comprised 36.181 decisions with 34.744 positive, 383 negative and 1054 neutral labels. For the evaluation, three meaningful attributes were selected, and the quality threshold t was set to 0.6. The evaluation run yielded 33 results that indicated a false "keep relation" label. Manually investigating these findings together with our expert partners in the two companies, we were able to verify the correctness for 30 of them. This equates to an accordance of 90,9%. Please note that a higher quality threshold would have resulted in more findings and presumably also a lower accordance, while a lower threshold would have had the opposite effect. Using the selected strict parameterization and the low generated number of suspicious reviewer decisions allowed us to conduct focused review workshops with the industry partners without confronting them with a potentially large number of results (which in turn would have complicated the expert review process). Summing up, the chosen parameterization allowed us to yield a small number of likely correct results with an acceptably small error rate. As an example, Fig. 2 displays four findings discovered by the algorithm regarding two employees in an access grid. The employees are displayed in the rows of the visualization. Four positive decisions (see violet border) were made by the same reviewer, while

⁷ Note that determining a suitable threshold must be done manually and the value highly depends on the imported data.

14 S. Groll et al.

all other permissions of the two employees were marked for removal by another reviewer. Expert consultations revealed that the reviewer who made the negative decisions was the manager of these employees and knew that they had left the company during the past month. At the same time the positive decisions were made by a system administrator who was not informed about this fact. According to our industry partner’s access review campaign experts, these findings are especially relevant, as they do not only expose possibly wrong decisions, but rather reveal structural and organizational problems concerning the quality of their overall IAM processes, e.g. lack in knowledge management or usability. Further expert evaluation revealed that reviewers had allegedly pressed a wrong button or decided too hastily.

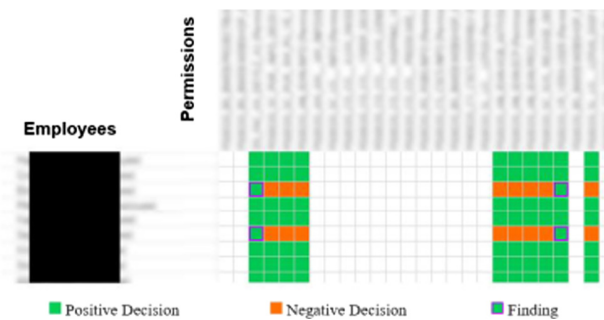


Fig. 2. Findings of the algorithm visualized with an access grid.

6 Conclusion

In the course of this work, we addressed the quality evaluation of access review decisions by answering the research question “How can low-quality review decisions in an access review be identified ex-post?”. To achieve this, we firstly developed a generalized problem formulation. Subsequently, we mapped this problem and its characteristics to the research area of crowd sourcing. We facilitated an approach by Geva and Saar-Tsechansky [16] developed to assess the quality of crowd labelling decisions and consequently adapted it to the given problem formulation. Additionally, we developed a theory-ingrained artifact that allowed us to identify erroneous access review decisions automatically. During result evaluation we were able to prove the applicability of our approach for evaluating the quality of human access review decisions. We discussed interim results and considered the feedback provided by our industry partners in order to make sure that the research objectives are in line with actual practical needs. Following the ADR principle of generalized outcome, we provided both, a generalized problem, as well as a generalized solution formulation. The approach may be applied back to the original approach from the crowd labelling realm which, in turn, leads to the achievement of a general improvement and contribution to the field.

References

1. The European Parliament and the Council of the European Union: General Data Protection Regulation. <https://eurlex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32016R0679>. Accessed 22 Mar 2021
2. One Hundred Seventh Congress of the United States of America: “Sarabanes-Oxley-Act (SOX)”. <https://www.iso.org/standard/54533.html> (2002). Accessed 22 Mar 2021
3. Basel committee on banking supervision: Basel III: a global regulatory framework for more resilient banks and banking systems. <https://www.bis.org/publ/bcbs189.pdf> (2010). Accessed 22 Mar 2021
4. International organization for standardization. “ISO/IEC 27002: Information technology - security techniques - code of practice for information security controls”. <https://www.iso.org/standard/54533.html> (2013). Accessed 22 Mar 2021
5. Federal office for information security (BSI): “IT-Grundschutz”. https://www.bsi.bund.de/EN/Topics/ITGrundschutz/itgrundschutz_node.html (2018). Accessed 22 Mar 2021
6. Azhar, A., Amin, M., Nauman, M., Shah, S.U.: Efficient selection of access control systems through multi criteria analytical hierarchy process. In: 2012 International Conference on Emerging Technologies, pp. 1–8 (2012). <https://doi.org/10.1109/ICET.2012.6375419>
7. Bobba, R., Gavrila, S., Gligor, V., Khurana, H., Koleva, R.: Administering access control in dynamic coalitions. In: Proceedings of the 19th Conference on Large Installation System Administration Conference, vol. 19, p. 23, LISA 2005. USENIX Association, USA (2005)
8. Brodley, C.E., Friedl, M.A.: Identifying mislabeled training data. *J. Artif. Int. Res.* **11**(1), 131–167 (1999)
9. Cole, R., Puro, S., Rossi, M., Sein, M.: Being proactive: where action research meets design research. In: ICIS 2005 Proceedings, p. 27 (2005)
10. Dekel, O., Shamir, O.: Vox populi: Collecting high-quality labels from a crowd. In: COLT (2009)
11. Erickson, L.B., Trauth, E.M., Petrick, I.: Getting inside your employees’ heads: navigating barriers to internal-crowdsourcing for product and service innovation (2012)
12. Ferraiolo, D., Kuhn, D.R., Chandramouli, R.: Role-Based Access Control. Artech House, Boston (2003)
13. Fuchs, L., Pernul, G.: Supporting compliant and secure user handling—a structured approach for in-house identity management. In: The Second International Conference on Availability, Reliability and Security (ARES 2007), pp. 374–384. IEEE (2007)
14. Fuchs, L., Pernul, G.: HyDRo – hybrid development of roles. In: Sekar, R., Pujari, A.K. (eds.) ICISS 2008. LNCS, vol. 5352, pp. 287–302. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-89862-7_24
15. Fuchs, L., Preis, A.: BusiROLE: a model for integrating business roles into identity management. In: Furnell, S., Katsikas, S.K., Liroy, A. (eds.) TrustBus 2008. LNCS, vol. 5185, pp. 128–138. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-85735-8_13
16. Geva, T., Saar-Tsechansky, M.: Who’s a good decision maker? data-driven expert worker ranking under unobservable quality (2016)

16 S. Groll et al.

17. Hill, L.: How automated access verification can help organizations demonstrate HIPAA compliance: a case study. *J. Healthcare Inf. Manag.* **20**(2), 116 (2006)
18. Hummer, M., Kunz, M., Netter, M., Fuchs, L., Pernul, G.: Adaptive identity and access management-contextual data based policies. *EURASIP J. Inf. Secur.* **2016**(1), 1–16 (2016)
19. Ihl, A., Strunk, K.S., Fiedler, M.: The influence of utilitarian and hedonic motivation on success in crowd work (2018)
20. Jaferian, P., Rashtian, H., Beznosov, K.: To authorize or not authorize: helping users review access policies in organizations. In: 10th Symposium on Usable Privacy and Security ({SOUPS} 2014), pp. 301–320 (2014)
21. Jin, X., Krishnan, R., Sandhu, R.: A unified attribute-based access control model covering DAC, MAC and RBAC. In: Cuppens-Boulahia, N., Cuppens, F., Garcia-Alfaro, J. (eds.) *DBSec 2012*. LNCS, vol. 7371, pp. 41–55. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-31540-4_4
22. Khattak, F.K., Salleb-Aouissi, A.: Robust crowd labeling using little expertise. In: Fürnkranz, J., Hüllermeier, E., Higuchi, T. (eds.) *DS 2013*. LNCS (LNAI), vol. 8140, pp. 94–109. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40897-7_7
23. Kittur, A., Chi, E.H., Suh, B.: Crowdsourcing user studies with mechanical Turk. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 453–456 (2008)
24. Kittur, A., et al.: The future of crowd work. In: *Proceedings of the 2013 Conference on Computer Supported Cooperative Work*, pp. 1301–1318 (2013)
25. Kuhn, D.R., Coyne, E.J., Weil, T.R.: Adding attributes to role-based access control. *Computer* **43**(6), 79–81 (2010)
26. Kunz, M., Puchta, A., Groll, S., Fuchs, L., Pernul, G.: Attribute quality management for dynamic identity and access management. *J. Inf. Secur. Appl.* **44**, 64–79 (2019)
27. Leicht, N., Rhyn, M., Hansbauer, G.: Can Laymen outperform experts? The effects of user expertise and task design in crowdsourced software testing (2016)
28. Meier, S., Fuchs, L., Pernul, G.: Managing the access grid-a process view to minimize insider misuse risks (2013)
29. Menon, A., Mehrotra, K., Mohan, C.K., Ranka, S.: Characterization of a class of sigmoid functions with applications to neural networks. *Neural Netw.* **9**(5), 819–835 (1996)
30. Oleson, D., Sorokin, A., Laughlin, G., Hester, V., Le, J., Biewald, L.: Programmatic gold: Targeted and scalable quality assurance in crowdsourcing. In: *Workshops at the Twenty-Fifth AAAI Conference on Artificial Intelligence*. Citeseer (2011)
31. Pfitzmann, A., Hansen, M.: A terminology for talking about privacy by data minimization: anonymity, unlinkability, undetectability, unobservability, pseudonymity, and identity management (2010)
32. Rouse, A.C.: A preliminary taxonomy of crowdsourcing (2010)
33. Sandhu, R.S., Coynek, E.J., Feinstein, H.L., Youman, C.E.: Role-based access control models. *Computer* **29**(2), 38–47 (1996)
34. Sein, M.K., Henfridsson, O., Purao, S., Rossi, M., Lindgren, R.: Action design research. *MIS Q.* **35**(1), 37–56 (2011)
35. Tavanapour, N., Bittner, E.A.: The collaboration of crowd workers (2018)
36. Valecha, R., Kashyap, M., Rajeev, S., Rao, R., Upadhyaya, S.: An activity theory approach to specification of access control policies in transitive health workflows (2014)

Monitoring Access Reviews by Crowd Labelling 17

37. Whitehill, J., Wu, T.F., Bergsma, J., Movellan, J., Ruvolo, P.: Whose vote should count more: optimal integration of labels from labelers of unknown expertise. *Adv. Neural Inf. Process. Syst.* **22**, 2035–2043 (2009)
38. Wöhner, T., Köhler, S., Peters, R.: Good authors = good articles?-how Wikis work. In: *Wirtschaftsinformatik*, pp. 872–886 (2015)

7 Digital Nudges for Access Reviews: Guiding Deciders to Revoke Excessive Authorizations

Current status:	Published
Conference:	20th Symposium on Usable Privacy and Security (SOUPS 2024), Philadelphia, August 11-13, 2024
Date of acceptance:	May 06, 2024
Full citation:	Thomas Baumer, Tobias Reittinger, Sascha Kern, and Günther Pernul. Digital nudges for access reviews: guiding deciders to revoke excessive authorizations. In <i>Twentieth Symposium on Usable Privacy and Security</i> , pages 239–258, 2024.
Authors contributions:	Thomas Baumer 40% Tobias Reittinger 25% Sascha Kern 25% Günther Pernul 10%

Conference Description: The 2024 Symposium on Usable Privacy and Security (SOUPS) will bring together an interdisciplinary group of researchers and practitioners in human computer interaction, security, and privacy. The program will feature: Technical papers, including replication papers and systematization of knowledge papers. Workshops and tutorials. A poster session. Lightning talks.

Digital Nudges for Access Reviews: Guiding Deciders to Revoke Excessive Authorizations

Thomas Baumer
Nexis GmbH

Tobias Reitinger
University of Regensburg

Sascha Kern
Nexis GmbH

Günther Pernul
University of Regensburg

Abstract

Organizations tend to over-authorize their members, ensuring smooth operations. However, these excessive authorizations offer a substantial attack surface and are the reason regulative authorities demand periodic checks of their authorizations. Thus, organizations conduct time-consuming and costly access reviews to verify these authorizations by human decision-makers. Still, these deciders only marginally revoke authorizations due to the poor usability of access reviews. In this work, we apply digital nudges to guide human deciders during access reviews to tackle this issue and improve security. In detail, we formalize the access review problem, interview experts ($n = 10$) to identify several nudges helpful for access reviews, and conduct a user study ($n = 102$) for the *Choice Defaults Nudge*. We show significant behavior changes in revoking authorizations. We also achieve time savings and less stress. However, we also found that improving the overall quality requires more advanced means. Finally, we discuss design implications for access reviews with digital nudges.

1 Introduction

The Open Web Application Security Project (OWASP) lists “broken access control” as the Top 1 vulnerability and discovers it in 94% of the tested web applications [36]. Excessive authorizations are one driver for this OWASP vulnerability, as these are granted without an actual need and thus open an unnecessary attack surface. More precisely and within this paper, we ask highly qualified Identity and Access Management (IAM) experts to estimate the ratio of excessive autho-

rizations in mid- and large-sized organizations. Our experts expect about a fifth to a quarter of the authorizations to be excessive and vulnerable ($M = 22.8\%$, $SD = 6.4\%$, $n = 10$).

To mitigate this vulnerability, regulative authorities demand organizations to evaluate their authorizations with periodic access reviews. Well-known regulations include SOX [52], Basel III [6], MARisk [12], or HIPAA [51]. In large organizations, this involves hundreds of access review deciders for six figures of authorizations [18, 39]. These deciders (e.g., department heads) evaluate these authorizations within their responsibility. Although accountable, deciders face a time-consuming and frustrating task, as their expertise and objectives might not primarily match with security. Responsible deciders must also avoid mistakes: While revoked authorizations can interrupt their organization shortly, falsely confirmed excessive authorizations drive security risks [25]. Research [18] shows in a real-world case study that deciders only revoke 1.2% of the reviewed authorizations instead of the expected one-fifth excessive ones. Besides this clear need for improvement, only a few papers [18, 22, 26, 39] study access reviews.

As shown by Jaferian et al. [26], crucial issues for access reviews are rooted in poor usability. Using digital nudges to guide decisions [53] is thus a promising approach to improve access reviews. However, we identify several research gaps: First, current research does not formalize access reviews. Second, it is unknown how digital nudges address access review challenges. Third, it is unclear whether digital nudges actually improve access reviews. We investigate these research gaps with the following research questions:

- Q1 *How to formalize the access review problem?*
- Q2 *How do access review challenges map to digital nudges?*
- Q3 *Does an applied digital nudge (the Choice Defaults Nudge) benefit the access review problem?*

This work follows a mixed methods approach in an exploratory sequential design. We use qualitative methods to define a formal and precise notation of the underlying problem (Q1) and to interview highly qualified experts ($n = 10$) about

Copyright is held by the author/owner. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee.

USENIX Symposium on Usable Privacy and Security (SOUPS) 2024.
August 11–13, 2024, Philadelphia, PA, United States.

applying digital nudges for access reviews (Q2). Moreover, our quantitative methods use insights of Q1 and Q2 to conduct a user study ($n = 102$) with an application of the *Choice Defaults Nudge* for access reviews (Q3). Consequently, our methods lead to the following contributions:

- We are first to formalize the access review problem.
- We map the expected effects of digital nudges to access review challenges based on 10 expert interviews. We find that access reviews benefit from digital nudges.
- We show behavior changes leading to quality improvements and more revoked authorizations by applying the *Choice Defaults Nudge* within a user study ($n = 102$). Moreover, we achieve time savings and lower frustration.

The remainder of this work is outlined as follows: Section 2 covers the background of this work, including relevant terminology, access review challenges, digital nudges, and related work. Section 3 provides details about our mixed method approach. Subsequently, we present the three-fold results of our paper. In Section 4, we first formalize the access review problem. Second, we map digital nudges with the access review challenges through the expert interviews in Section 5. Third, we conduct a user study on the *Choice Defaults Nudge* for access reviews in Section 6. Following the results, we discuss the general findings of this work in Section 7. Finally, Section 8 concludes and gives an overview for future work.

2 Background

2.1 Terminology

Identity and Access Management (IAM) is a cornerstone of modern cybersecurity, as it manages users and their access to sensitive data and services of organizations. Therefore, IAM provides tools to administer, authorize, and authenticate identities. Regulative authorities acknowledge the relevance of IAM and demand proper security controls. Besides state-of-the-art authentication, one crucial control is to demonstrate that the users still require granted access. Access reviews are a typical tool to prove the actuality of the granted access. These access reviews are the main focus of this work.

Access reviews are a periodic and compliance-driven process to verify users' authorizations. A team of domain experts, managers, application owners, and security admins typically reviews the granted authorizations with their knowledge of current processes, people, and resources. Especially in large organizations, access reviews are labor-intensive. Because of the recurring workload of access reviews, an organization might not finish an access review before the next one starts. The primary goal of access reviews is *revoking excessive authorizations*. Secondary goals are the determination of responsibilities for authorizations, requesting missing authorizations, or organization-specific data quality requirements. [18,22,26]

Nudges help humans make choices in analogous and digital systems. While these individuals must make their choices freely, *choice architects* design *choice architectures* to support their decisions by *nudging* towards a desired option. A nudge is thus a characteristic, influencing a decision in the interest of the decider. An example of a nudge in a supermarket is making healthy food easily accessible while making the unhealthy one harder to reach. From an ethical perspective, a nudge does not prevent a human from making a specific choice and only influences the decision in the best interest of the human. A *digital nudge* applies the idea of nudges to information systems. With features of user interfaces for guidance, users can make their choices freely and supported by the best advice of the choice architecture. [27,45,50]

2.2 Access Review Challenges

Based on expert interviews, Jaferian et al. [26] summarize access review challenges (C1-C5). We utilize these challenges throughout the paper, and thus detail them in this section.

C1: Scale outlines the number of involved IAM entities for the access review. The scale of the users, roles, permission, accounts, or assignments quickly grows into large numbers [18,26,39], making careful considerations for organizing the access review's workload necessary. Furthermore, the heterogeneity of these entities within real-world organizations intensify this challenge [30].

C2: Lack of Knowledge refers to the understandability of roles and permissions [26,30,31]. IAM entities might not have telling names, comprehensive descriptions, or concepts like roles or permissions might not have been fully understood. Experts thus might take uninformed or *best guess* decisions, leading to a bias for keeping unnecessary granted authorizations, violating the Principle of Least Privilege (PoLP) [18]. Additionally, for large organizations, the knowledge about these entities is distributed (or even missing completely), making the advice of responsible domain experts necessary [30].

C3: Frequency describes a dilemma for the managers: access reviews are not their *actual* responsibility, but they are frequently asked for it [26]. The experts might not feel a need to participate, leading to failing access reviews. Ultimately, this may cause even more access reviews, since successfully executed access reviews are part of compliance and audits. Thus, while access reviews usually are only required yearly, some organizations execute them quarterly, hoping not to fail access reviews due to lack of participation [26].

C4: Human Errors are common due to the scale and manual execution of access reviews by human deciders. These experts ultimately decide about required or excessive access by applying the best of their knowledge. This process is, therefore, inherently error-prone, as decisions to the best of the experts' knowledge might be incorrect or uninformed [18,26].

C5: Exceptional Cases occur due to the scale and complexity of access reviews. Context knowledge is sometimes

required for an informed authorization decision. For example, some members of organizations might replace others while on leave, trainings or tests might require temporary access, etc. might cause disturbances [26].

2.3 Digital Nudges

Based on a literature survey, Jesse and Jannach [27] propose a taxonomy for digital nudges. The authors distinguish four primary categories with further sub-categories of digital nudges (N01-N13): decision information (N01-N04), decision structure (N05-N08), decision assistance (N09-N10), and social decision appeal (N11-N13). We refer to these nudges throughout the paper, and thus explain them in this section.

Decision Information tries to present information helpful for the decision-maker without altering the available choices. This category comprises information translation (N01), salience (N02), visibility (N03), and phrasing (N04).

- *N01: Information Translation* targets reducing the cognitive effort for a decision by simplifying information or decreasing vagueness and ambiguity [48].
- *N02: Information Salience* aims to raise or decrease the prominence of information, by visualizations or making information harder or easier to notice [11, 48].
- *N03: Information Visibility* fosters decision information. This category includes mechanisms to disclose [24, 28, 48], compare [11, 48] or warn with [24, 33, 48] (tailored [28, 33] or external [35, 49]) information.
- *N04: Information Phrasing* puts presented information in context to intervene with the decisions to make. This category comprises the utilization of heuristics or biases like anchoring [33, 34, 48], availability [44, 50], the endowment effect [11, 44], framing [11, 33, 48], loss aversion [34, 44, 48], priming [11, 48, 50], etc.

Decision Structure alters the decision arrangement, comprising the decisions' range & composition (N05), defaults (N06), consequences (N07), and required effort (N08).

- *N05: Range & Composition* groups and categorizes choices. Therefore, choice architects or the decision-makers themselves break large decision structures into smaller category partitions [28, 48, 53], to present these one after another [28, 33] or to make them more comparable to each other [28, 35]. Choice architects can also utilize ordering effects for the presented options [11, 48].
- *N06: Choice Defaults* is one of the most effective and well-studied nudges [24]. The nudge preselects choices without hindering the decision maker from actively making another choice. On the one hand, a decision-maker is more invested in an actively made decision [35, 48]. On the other hand, decision-makers

rather accept the preselected status quo than actively decide against it [11, 24, 35, 48].

- *N07: Option Consequences* add further yet rational insignificant effects to the choice without changing the overall economic incentives. These consequences include social outcomes or minor benefits & costs [24, 35].
- *N08: Option-related Effort* modifies the effort or ease to make decisions. This nudge includes capping [11, 48] or raising financial & physical effort [24, 35] of decision-makers choices to mitigate mindless actions. Furthermore, eased and more convenient choices speed up decisions, e.g., making desired choices more accessible [48].

Decision Assistance aids decision-makers to realize their intentions. This category includes the usage of reminders (N09) and commitment facilitation (N10).

- *N09: Reminders* actively put already available information into or out of the attention focus of the decision-makers. This nudge includes reminding of underlying goals, deadlines, and their relevance [11, 24, 33, 35, 48, 49] or stating social expectations for decisions [35].
- *N10: Commitment Facilitation* helps decision-makers to (timely) finish their asked for decisions. This nudge includes precommitment strategies (e.g., user-defined sub-goals) [24, 33, 35, 48] or public commitment (e.g., pressure by publicly communicating own goals) [11, 35].

Social Decision Appeal category focuses on the social implications of nudges, including the Messenger Reputation (N11), Social Reference Point (N12), and Empathy Instigation (N13).

- *N11: Messenger Reputation* considers the reputation of the messenger delivering the information for the nudge. On the one hand, the messenger effect nudges a decision-maker since the messenger provides a certain and influencing impression about itself. For example, an actually well-designed and important choice architecture might dilute its seriousness if it contains many spelling mistakes [44]. On the other hand, the reputation of a system can be improved when choice architects expect and forgive the errors of their decision-makers [28, 53].
- *N12: Social Reference Point* nudges a decision based on social opinions. E.g., the opinion of a majority (Argentum-Ad Populum), group (Group-Ad Populum) [16], or an opinion leader [35] can influence decision-makers. Additionally, deciders tend to follow a herd [34, 44, 48] and might desire a comparison with their peers influencing their own decisions [24, 33, 35].
- *N13: Empathy Instigation* uses feelings to influence deciders. For example, an avatar might smile or cry upon the choices of a decision-maker (moral suasion) [11, 48], or a choice architect can trigger reciprocity by doing something good for the decision-makers to nudge them into returning the favor with good choices [11].

2.4 Related Work

Access control ensures users can only act within their intended authorizations and is characterized by its necessary yet cumbersome maintenance. Related research on maintenance covers more efficient access control models, optimization, and general maintenance processes like access reviews. By evolving from access control matrices [41], the most dominant access control models are Role-Based Access Control (RBAC) [15, 37, 43] and Attribute-Based Access Control (ABAC) [23, 46] as these reduce maintenance costs. Modeling access control policies considers bottom-up, top-down, or hybrid approaches [14] but often overlook their actual optimization without recalculating them [31, 38]. Therefore, access control maintenance targets up-keeping authorizations in changing needs and environments based on IAM goals [25, 30]. This includes periodically reviewing and revoking excessive access [18, 22, 26], granting missing access [47, 54], and timely propagation [7] to maintain secure authorizations. This paper especially relates to work on maintenance by access reviews: Jaferian et al. [26] study its challenges and usability. Puchta et al. [39] show positive effects on using external data for access reviews. Groll et al. [18] assess decision quality. Hill [22] conducts a case study for HIPAA [51] compliant access reviews.

Digital nudges are a popular research topic, as shown by various surveys: While Bergram et al. [9] conduct a general literature review, Schaer and Stanoevska-Slabeva [44] analyze digital nudges in customer-journeys and Jesse and Jannach [27] in recommender systems. Additionally, a survey of Caraban et al. [11] covers a practical and ethical application. As an established means to shape human behavior, applications of (digital) nudges exist for many domains. Examples include e-commerce [2, 13], sustainable smart home [8], contract tracing [17], or cybersecurity. In detail, cybersecurity examples include digital nudges to prevent phishing [55] or increase password quality [29, 56, 57]. An application of digital nudges for access reviews remains open so far.

3 Methods

This research uses mixed methods in an exploratory sequential design. First, we use qualitative methods to formalize the Access Review Problem (ARP) (Q1) and relate access review challenges to digital nudges (Q2). Second, we use these qualitative insights in quantitative methods to study the effect of the *N06: Choice Defaults* for access reviews (Q3). Third, a discussion wraps up the findings. Figure 1 depicts our mixed methods. In the following, we detail each part.

3.1 Formalizing the Access Review Problem

While the access review challenges comprise a global view, we formalize the actual Access Review Problem (ARP) in

Section 4. Its goal is to understand the underlying problem better. This formalization targets a quantifiable and comparable foundation for the solution of the ARP. Thus, we argue access review as a transition between two authorization states, depicted as confusion matrices. This precise formalization of the ARP is the basis for the hypotheses of the user study.

3.2 Relation of Access Review Challenges to Beneficial Digital Nudges

Complementary challenges to the ARP are discussed in the literature, including scale, lack of knowledge, frequency, human errors, and exceptional cases [26]. Digital nudges are a promising approach to address the ARP and its challenges. But it is unknown, whether digital nudges can help and which effects can be expected from their application (Q2).

To better understand this relationship between access review challenges and digital nudges, we investigate and map access review challenges from Jaferian et al. [26] with the digital nudge taxonomy of Jesse and Jannach [27] by conducting semi-structured expert interviews based on the guidelines of Adams [1]. The interviewed industry experts provide practical experience in access control and reviews. Therefore, we target highly qualified professionals with at least five years of experience working with large IAM systems, periodical executed access reviews, and managing thousands of identities or consultants with practical experience for many enterprises. Of course, these highly qualified experts are not readily available, but we managed to acquire 10 of these experts through personal and professional contacts. The experts are located in Germany. We use their expertise for a well-grounded argumentation for the relationship between access review challenges and digital nudges. Section 5.1 details further on the method for the expert interviews.

3.3 User Study for the Choice Defaults Nudge

After laying out theoretical foundations for digital nudges and access reviews in Sections 4 and 5, we study the application of a selected digital nudge in-depth. The expert mapping of digital nudges and access review challenges suggests several digital nudges. To sharpen the scope of the use study, we select *N06: Choice Defaults* based on the following reasons:

- Literature considers *N06: Choice Defaults* among the most effective digital nudges [24].
- The expert interviews had strong positive and negative expectations, inviting a more detailed examination.
- We felt confident to apply the *N06: Choice Defaults* to an access review and study its effects precisely.

We design an access review, simulating a real case: experts often describe access reviews as repetitive, time-consuming, and tedious tasks, requiring a strenuous thought process to




Figure 1: Mixed methods approach for this study.

determine correct authorizations. We thus hand out each participant a one-pager about the case. Participants manage a fictitious marketing department containing three teams within the case: graphic design, social media, and event management. The instruction describes the functions and tasks of each team and explains the unwanted implications of excessive or missing authorizations. While it is theoretically possible to review each decision using the document correctly, it takes some thought to make a correct decision.

To study the *N06: Choice Defaults* in-depth, we use three distinct configurations for access reviews with the same data basis: default accept, default reject, and a neutral default. This directly compares the default accept and reject configuration with a neutral state. The default accept configuration preselects every decision with an accept, the default reject vice versa, and the neutral default does not preselect.

We acquire 102 participants from a university context in Germany and randomly assign them to one of the three *N06: Choice Defaults* configurations. We select our sample size based on similar papers (c.f. Caine [10], also for the expert interviews). The (under-)graduate students have mostly a background in business informatics and IT security, indicating that they know essential IT security concepts and enterprise information systems. Furthermore, the participants are unaware of the research objective on digital nudges. We raffle a €100 gift card to one lucky participant to motivate participation. The participants must log in with authenticated accounts to avoid repeated participation and enable remote participation. We pilot the study with fellow researchers. Section 6.1 provides further details for the method.

3.4 Ethical considerations

Our experts were informed and consented to an anonymous publication of parts of their interviews. We will not share the recordings and delete the data one year after the publication.

The Institutional Review Board (IRB) *German Association for Experimental Economic Research e.V* approved the user study to comply with ethical requirements for working with humans. The certificate is available online.¹

¹<https://gfew.de/en-ethik/HQwmKGTZ>

4 Q1: The Access Review Problem

To better understand access review and benchmark our user study design, we introduce a representation of granted authorizations and security policies within a confusion matrix depicting User Permission Assignments (UPAs). Figure 2 maps the actually granted authorizations with security policies. We assess authorizations as effective access grants (which may contain errors), while security policies define the conceptual access users should have (ground truth). We construct a classical confusion matrix by mapping these authorizations and security policies with a binary distinction. Thus, the effectively granted UPAs are Predicted Positive (PP) as $PP = TP + FP$, while $P = TP + FN$ should be granted. PN and N are vice versa not-granted UPAs. Therefore, the True Positives (TPs) describe UPAs, granted in reality and conceptually. The sensitivity $SEN = \frac{TP}{P}$ represents the rate of correctly granted UPAs. Vice versa, True Negatives (TNs) describe UPAs, not granted in reality and in concept. The specificity $SPC = \frac{TN}{N}$ represents the rate of correctly not-granted UPAs. Together, sensitivity and specificity express the balanced accuracy $BA = \frac{SEN+SPC}{2}$, equaling 100% in a perfect world without errors.

		Authorization	
		Positive PP	Negative PN
Security Policy	Positive P	TP	FN
	Negative N	FP	TN

Figure 2: Confusion matrix for UPAs.

However, type I (False Positives (FPs)) and type II (False Negatives (FNs)) errors are present in reality. On the one hand, FPs are granted authorizations not considered by security policies (excessive UPAs). These excessive authorizations drive security risks, as over-privileged users are a target for threat actors. The primary goal for access reviews is lowering FP, which is highlighted in Figure 2. On the other hand, FNs are mistakenly not granted authorizations (missing UPAs). An example of their impact is when users cannot do their legitimate tasks because they do not have access to the required systems. This causes dissatisfaction for the users and slows down processes. In a relative notation, the False Discovery Rate (FDR) describes the percentage of excessive UPAs FP

based on PP as $FDR = \frac{FP}{PP}$. Vice versa, the False Omission Rate (FOR) describes missing UPAs as $FOR = \frac{FN}{PN}$.

Thus, an access review can be understood as transitioning from one UPA set depicted as confusion matrix C^1 to another C^2 . The primary goal is to reduce the FDR while retaining or improving BA. We introduce definitions for Access Reviews (ARs) and the Access Review Problem (ARP) as:

Definition 4.1 (Access Review (AR)). Given a confusion matrix C^1 describing an UPA^1 set, an access review AR revokes a subset of the effectively granted authorizations $R \subset PP^1$. When executing AR a confusion matrix C^2 describes the resulting set as $UPA^2 = UPA^1 \setminus R$.

Definition 4.2 (Access Review Problem (ARP)). Design AR in such a way that a (human) deciders can review and revoke UPAs $R \subset PP^1$ according to their knowledge about security policies P^1 , that the FDR is reduced ($FDR^1 > FDR^2$), without lowering BA ($BA^1 \leq BA^2$). The ARP is solved on a $FDR^2 = 0\%$ without decreasing BA : $BA^1 \leq BA^2$.

The following hypotheses hence allow testing whether an access review design improves the ARP:

- H₀** An access review design does not improve the ARP as the FDR remains or rises $FDR^1 \leq FDR^2$ or BA remains or decreases $BA^1 \geq BA^2$.
- H₁** An access review design improves the ARP as the FDR decreases $FDR^1 > FDR^2$ and the BA raises $BA^1 < BA^2$.

5 Q2: IAM Experts on Digital Nudges

5.1 Method Details

The interviews comprise three phases: an interviewee introduction, an explanation of access review challenges and digital nudges, and a workshop to generate the mapping of access review challenges and digital nudges. (i) The interviewee's introduction collects data about their access review experience, their perspective on its challenges, and their estimation of excessive authorizations (FP). (ii) The explanation phase ensures essential knowledge about digital nudges, reminds the interviewee of access review challenges, and ensures a common vocabulary. We use the interviewees' perspectives on access review challenges to explain to them the access review challenges of Jaferian et al. [26]. (iii) The procedure for querying the mapping for each considered digital nudge [27] follows this scheme: First, we explain the digital nudge in general and provide a suitable example for the interviewee. Afterward, we let the expert freely reflect on the effect of this digital nudge and its benefit to all access review challenges. Finally, we ask the expert to rate each access review challenge on a five-level Likert scale from very positive (+2) to very negative (-2). This rating scheme helps the expert to express

their arguments more comparable to each other. The complete interview script is available in Appendix A.1.

We interviewed 10 highly qualified experts with experience in conducting several Access Reviews (ARs) specialized for IAM by implementing IAM tools (engineers), responsible for managing thousands of users in IAM systems (inhouse), or advising clients (consultants). Table 1 protects their identities but depicts their high expertise for ARs. The interviews took an average of 60 minutes and were recorded, transcribed, coded, and evaluated. We translated relevant parts of the interviews into English during the coding process.

Table 1: Participants for expert interviews.

Interview	Experience				Sector
	Years	Clients	Users	ARs	
E01: IAM consultant	8	40		20	Multiple
E02: IAM consultant	5	15		10	Multiple
E03: IAM engineer	12	40		15	Multiple
E04: IAM inhouse	8	15	1k	40	Insurance
E05: IAM consultant	19	25		10	Multiple
E06: IAM consultant	13	40		25	Multiple
E07: IAM consultant	6	15		50	Multiple
E08: IAM inhouse	15	2	19k	4	Biotech
E09: IAM consultant	11	4		10	Banking
E10: IAM inhouse	7	1	13k	120	Insurance

We recorded the interviews with Microsoft Teams, transcribed them with Word, and summarized and coded them with Excel. For the coding, we use both deductive and inductive coding [3]. Since we already know the access review challenges [26], we first applied deductive coding based on these challenges for each digital nudge. This deductive coding already sorts large parts of the interviews in proven codes. However, we noticed that several augmentations exist within these codes. Thus, we also developed inductive codes for each nudge and challenge combination to capture the interviews comprehensively. For the rating of each nudge and challenge pair, we initially used the mean expert ratings. After coding and comparing the interviews, we slightly adapted the ratings, to balance well-reasoned arguments across the experts. The resulting codebook is available in Appendix A.2.

5.2 Results

This section presents the experts' mapping. We build on the presented background of the access review challenges (C1-C5) in Section 2.2 and digital nudges (N01-N13) in Section 2.3. The resulting mapping is depicted in Table 2, whereas the challenges serve as columns and the digital nudges as rows. The cells summarized a rating for each challenge and nudge. In the following, we detail each digital nudge.

N01: The experts stress the benefits of comprehensible data. While C1 and C3 do not decrease, comprehensible data indirectly increases its learnability and comfort for the deciders, easing management eventually. For C2 and C4, the

Table 2: Nudges [27] and access review challenges [26].

Nudges	C1	C2	C3	C4	C5
N01: Information Translation	1	2	1	2	0
N02: Information Salience	1	0	1	1	2
N03: Information Visibility	1	2	0	1	2
N04: Information Phrasing	0	-1	0	1	0
N05: Range & Composition	2	1	1	2	2
N06: Choice Defaults	2	-2	2	-2	0
N07: Option Consequences	0	-1	1	-1	-1
N08: Option-related Effort ↗	-1	1	-1	1	1
N08: Option-related Effort ↘	1	-1	1	-1	-1
N09: Reminders	0	1	2	-1	0
N10: Commitment Facilitation	1	0	1	1	0
N11: Messenger Reputation	1	2	1	2	2
N12: Social Reference Point	0	2	0	1	2
N13: Empathy Instigation	1	1	1	1	0

Note: Option-related effort is ↗ = increased, ↘ = decreased. The Likert scale spans from very positive +2 to very negative -2.

experts anticipate a strong positive effect, as comprehension is essential for C2: “If data is displayed more comprehensibly, it’s helpful for users with little knowledge [C2] about the decision.” (E06). Being comfortable with the data is relevant for C4: “If the user is comfortable with the displayed data, you can expect fewer human errors [C4].” (E07)

N02: The experts emphasize the focus: “In my opinion is the highlighting of C5 the only option to manage large data sets.” (E05) However, “it depends on the quality of the highlighting” (E03), since excessive or missing highlighting might draw away attention from relevant decisions. But upon sufficient and reliable quality, decision-makers can efficiently focus on the highlighted decisions or attributes and decide the remainder quicker (benefit for C1 and C3). Decision-makers “actually want to decide diligently but are hindered by its scale. These decision-makers could diligently and mindfully decide just the highlighted decisions in an efficiency tradeoff.” (E09)

N03: Showing additional data is crucial for C2 and C5 to make informed decisions while streamlining the focus to relevant attributes (C4). By only offering limited attributes for each decision in default, the management of C1 is eased. However, the user might not know the relevancy of specific hidden attributes as these move out of focus (C4).

N04: Our interview partners express reservations as decisions might not be based on rational knowledge but on biased phrasing (C2). However, for a well-executed implementation, its utilization can raise the access review acceptance (C4) as its relevancy could be communicated more effectively.

N05: The setup of meaningful partitions and sorting imposes overhead compared to just showing all decisions in one turn, thus worsening C1 and C3. However, the experts anticipate quite positive effects on all challenges. Similar sorted or clustered partitions leverage efficiencies as deci-

sions transfer to whole partitions. These efficiencies ease the management for C1 and C3 since the workload decreases, while more consistent and mindful decisions mitigate C2 and C4. Furthermore, clustering and appropriate communication of exceptional cases (C5) can positively influence.

N06: The experts discuss the strong effects of *N06: Choice Defaults*. Due to the reduced workload by the preselection, the experts rate a positive effect on C1 and C3. However, the experts worry that deciders adopt a preselected default without further thought, leading to uninformed (C2) and mindless (C4 and C5) decisions. While a mindful default prevents errors on uncertainty (like for C5) or on evident cases, just adopting the recommended default can become a fallacy, assuming the recommendation algorithm’s imperfections. This is especially an issue if the decision-makers trust the preselection so that they mindlessly adopt the default instead of a mindful decision. A falsely set default would then lead to a systematic bias, endangering the next audit relevant to compliance. In sum, the experts anticipate the potential of *N06: Choice Defaults*, but advise careful application.

N07: “In practice, negative consequences dominate. For example, we will tell your boss if you don’t finish your access review tasks within 14 days.” (E01) The experts acknowledge that creative and positive consequences could be feasible and reasonable, making frequent access reviews more comfortable (C3). However, they doubt there would be a game-changer in the long term because the effects would wear down over time (C3), and the decisions might be based on avoiding pressure or pursuing benefits (C4) instead of reason (C2 and C5). In this context, it is also worth noting that “disadvantaged individuals need special consideration” (E09) because finishing an access review in time might not be fair for these (C5).

N08: This nudge’s influence on the access review challenges is ambivalent, as it depends on whether the option-related effort is increased (↗) or decreased (↘). If the effort increases (vice versa for decrease), the users take more time to decide. For C1 and C3, this worsens the situation as the workload rises with its time consumption. Taking more time for a decision (e.g., requiring a reason for confirming a high-risk authorization) also benefits C2, C4, and C5, as the decider would need to consider a reason or reconsider the decision. But the experts also stressed the efficiency and acceptance of the access review, as some users easily become annoyed by increased effort: E.g., “We once required the users to set a note for the reviewed authorizations, but one user just put question marks for every note to bypass the input check.” (E04)

N09: “By a simple reminder [email], we observe more participation.” (E10) While reminders are especially relevant for C3 to communicate open tasks or instructions and goals for access reviews (C2), they can also pressure decision-makers to decide quickly but uninformed (C4). The audience and channel of reminders are also essential for C4. E.g., an inexperienced decider might require instructions or training. The experts also noted that reminders via an email channel

dominate in practice but are quickly perceived as spam. “Everybody wants something from all colleagues. Ironically, some colleagues even configure automated email filters which they won’t check afterward.” (E09) In this sense, a personal or multichannel address is most effective, but it is a considerable effort for the IAM team conducting the access review.

N10: The experts appreciate the autonomic commitment in combination with semantic partitioning (N05) of the decisions. An autonomic configuration of sub-goals and sub-deadlines suitable for the deciders benefits C1 and C3 as the deciders “perceive control over scale and frequency” (E06). This leads to more comfort, as sub-goals and sub-deadlines become meaningful for the deciders, mitigating C4.

N11: The experts stress the importance of this nudge: “Most important point; If the IAM team is not accepted, it is going be tough.” (E04) Furthermore, they note its failure in practice: “Access reviews are usually perceived negatively.” (E10) With a suitable messenger reputation, users will trust and endure the tedious tasks of the access reviews more, which is beneficial for C1 and C3. The experts also anticipate strong benefits for C2, C4, and C5 as the decision-makers will dare to ask or tell an approachable IAM team their relevant questions or mistakes: “If the IAM team is approachable, users communicate errors more eagerly or at all.” (E07)

N12: Similar to N11, if the social reference point sympathizes with the access review, decision-makers are likely to endure the tedious workload (C1 and C3). However, on low sympathy, the opposite effect might apply. The experts anticipate positive effects for C2, C4, and C5 because deciders discuss the access review: “For example, we introduced access review chat groups for business units. Decision-makers can talk about access reviews, like showing their own or seeing others’ progress, asking questions, etc.” (E07) In this sense, exceptional cases (C5) might become evident after a discussion and sharing knowledge about similar cases (C2), while noticing the colleagues’ progress might remind stragglers or expose them to peer pressure (C4).

N13: “On large scale [C1] and high frequency [C3], the decision-makers want to work with a pleasant tool.” (E06) Moral suasion and empathetic feedback (C2) can inform and convince the decision-maker about odd user behavior (e.g., mindlessly accepting all authorizations) without losing their motivation (C4). Reciprocity also fosters mitigation of C4 by “always addressing the positive side: the access review is meant to help you, the decision-maker, to compliantly and securely maintain your authorizations.” (E08)

Furthermore, the experts estimate a mean on excessive authorizations (FP) at 22.8% ($SD = 6.4\%$). Since we also asked our experts about common AR challenges, we confirm the AR challenges first published by Jaferian et al. [26].

In summary, our experts conclude positive and negative effects when using digital nudges. Table 2 summarizes these key takeaways. We hope to motivate future work with it as most digital nudges invite dedicated research on access reviews.

6 Q3: Choice Defaults in Access Reviews

6.1 Method Details

In the data set of the user study (Appendix B.1), we let participants review (accept or remove) granted UPAs $PP = TP + FP$ (legit TP or excessive FP), leading to UPA revoke operations only. Not granted UPAs $PN = FN + TN$ (missing FN or legit TN) are not considered. After piloting, we determined 160 UPAs serving as decisions to align an estimated study duration of 20-30 minutes and not to deter participation. Therefore, the crafted data set comprises 160 UPAs (PP) split into 80 legitimate ones (TP) and 80 excessive ones (FP), clearly distinguished by a case study document (see Appendix B.2). Figure 3 summarizes the initial UPAs as a confusion matrix.

		Authorization	
		$PP = 160$	$PN = 232$
Security Policy	$P = 80$	$TP = 80$	$FN = 0$
	$N = 312$	$FP = 80$	$TN = 232$

Figure 3: Confusion matrix for the case of the user study.

We configure and execute the access reviews with the commercial tool NEXIS4². The tool can import our data set, configure *N06: Choice Defaults*, execute large-scale access reviews, and collect relevant data points. Figure 4 displays a simplified screenshot of the review process. Further screenshots for all groups are available in Appendix B.3. For data collection, we make three observations for each access review participant: their decisions for the 160 UPAs, their time consumption, and their self-assessment for the NASA Task Load Index (TLX) [20]. (i) The tool stores each binary decision out-of-the-box, leading to a total of 16,320 manual decisions for 102 participants and 160 UPAs. (ii) We measure the time consumption for each participant by comparing the events for starting the access review and confirming the final completion prompt. (iii) After completion, we ask the participants to fill out a questionnaire for the NASA TLX [20] to capture their perceived workload. These questions are based on a Likert scale (-3 to +3) and include:³

- Mental Demand: How mentally demanding was the task?
- Temporal Demand: How hurried or rushed was the pace of the task?
- Performance: How successful were you in accomplishing what you were asked to do?
- Frustration Level: How insecure, discouraged, irritated, stressed, or annoyed were you?

²<https://nexis-secure.com/en/>

³We omitted the questions for physical demand and effort, as these are not applicable or relevant for our study.

	Employee	Permission
<input type="button" value="Approve"/>	<input type="button" value="Remove"/>	Moore, Evelyn F:\Documents\Social_Media_Strategy\
<input type="button" value="Approve"/>	<input type="button" value="Remove"/>	Moore, Evelyn Approval vacation requests
<input type="button" value="Approve"/>	<input type="button" value="Remove"/>	Miller, Sophia Book tradefair / exhibition stands

Figure 4: Simplified screenshot of the access review.

During the post-processing of the study, we used Microsoft Excel and R⁴ for data cleansing or data analysis. Data cleansing primarily comprises capping the time consumption for the AR to 60 minutes, as some participants took a break. We calculate the means, standard deviations, and non-parametric ANOVA of the AR confusion matrix, time-consumption, and NASA TLX indices. For our exploratory analysis of correlations (Spearman) and local regressions, we utilize a pair plot generated in R (see Appendix Figure 9). Supporting the open data idea, we publish all data to replicate our results on GitHub: <https://github.com/AccessReview/Availability>.

6.2 Results

This Section summarizes our observations of the user study (see Table 3). A post-hoc power analysis based on ANOVA for our three groups ($n = 34$) and an $\alpha = .05$ results in effect powers of .13 for a small effect ($f = .1$), .6 for a medium effect ($f = .25$), and .95 for a large effect ($f = .4$).

For all 102 participants, the mean review time t for the 160 decisions is $t = 22$ minutes with $SD = 13$ minutes. Deciders of all groups used to over-accept authorizations, amounting to a total accept rate of $1 - \frac{R}{PP} = 56.1\%$ (rather than a $SEN = 50\%$). H_0 is rejected for 99 of 102 reviews. The remaining 3 participants failed to achieve an ARP improvement. All participants' mean BA increased from 87.2% to 91.2% ($SD = 7.9\%$). The false discovery rate FDR , which represents the amount of excessive authorizations, was reduced from 50.0% to 21.6% ($SD = 14.7\%$). This improvement came at the cost of some erroneous revokes, leading to a mean FOR of 2.9% ($SD = 3.5\%$). In sum, most participants improved the ARP. The result data shows that two deciders behaved as “spammers” by either blindly accepting all authorizations (one decider in the accept group) or blindly rejecting them all (one decider in the reject group). These participants are among the three who failed to improve the ARP. While the data set is too small to make this finding statistically significant, it seems evident that the spammers just adopted the default.

The neutral configuration group is a control group for the default accept and reject nudge. Users from this group took a mean time of $t = 26$ minutes ($SD = 15$) and accepted 57.8% of the authorizations. The neutral group estimated the temporal demand as slightly low, with a mean score of $-.8$. On average,

⁴<https://www.r-project.org/>

neutral users stated the mental demand to be slightly high (.9) and their frustration to be neutral to slightly high (.5). They estimated their performance to be slightly above average (.9). The achieved BA is 91.9% ($SD = 5.8\%$), with the error rates FDR of 21.0% ($SD = 10.7\%$) and FOR of 2.6% ($SD = 2.5\%$).

The accept group only took $t = 19$ minutes ($SD=10$). With a time save of 24.3% to the neutral group. While the perceived TD was unchanged at $-.8$, both FL and MD were reduced by almost one point to a score of $-.2$ ($\Delta = -.7$) and $.2$ ($\Delta = -.7$). The accept rate was slightly higher than in the neutral group with 58.7% (+.9%). The default accept group achieved a BA of 92.3% ($SD = 5.3\%$), scoring .4% higher than the neutral one. The error rates were also marginally better than in the neutral group with $FDR = 20.8\%$ ($\Delta = -.2\%$, $SD = 9.3\%$) and $FOR = 2.2\%$ ($\Delta = -.4\%$, $SD = 2.6\%$).

Like the accept group, deciders of the reject group finished quicker than the neutral group with $t = 21$ minutes ($\Delta = -16\%$, $SD = 13$). Again, the estimated TD of $-.4$ did not reflect this ($\Delta = +.4$), but the stated FL and MD were reduced to $-.6$ ($\Delta = -1.1$) and $-.2$ ($\Delta = -1.1$). Unlike the accept group, however, the reject group showed a considerably reduced accept rate of 51.8% (-6.0%), which is very close to the initial $SEN = 50\%$. Unfortunately, the increased willingness to revoke did not improve the results: The deciders revoked fewer excessive authorizations than the neutral group ($FDR = 22.9\%$, $\Delta = +1.9\%$, $SD = 21.4\%$) and more correct ones ($FOR = 3.9\%$, $\Delta = +1.3\%$, $SD = 4.7\%$). With $BA = 89.4\%$ ($SD = 11.2\%$), BA was still improved regarding the initial state ($\Delta = +2.2\%$), but worse than the neutral configuration ($\Delta = -2.5\%$).

We ran a non-parametric Kruskal-Wallis test ($\alpha = .05$) to check for the significance of our observations between the three groups. We detect differences for the number of revokes R ($p = 0.039$), indicating that $N06: Choice Defaults$ did affect users' willingness to accept or reject authorizations. We also confirm differences for MD ($p = .049$) and FL ($p = .038$), indicating that lower stress perceptions result from the applied $N06: Choice Defaults$. We used Dunn's test for pairwise comparisons, showing that the neutral and reject groups differ for MD ($p.adjust = .045$) and FL ($p.adjust = .038$). However, the quality metrics BA , FDR , and FOR did not differ significantly between the study groups, which is unsurprising since the data set balances TP and FP at 80.

A test for Spearman correlation showed no significant correlation between review duration t and any of the quality metrics (BA , FDR , FOR), indicating that quality did not depend on the time spent. The data shows a significant positive correlation between the deciders' frustration level FL and t for the total population (.286) and the neutral group (.403), as well as between the stated mental demand MD and t (total: .237, neutral: .423). FL and MD are strongly correlated for all groups (total: .646, neutral: .589, accept: .672, reject: .664). We follow that deciders did not strictly distinguish between MD and FL and that longer reviews are perceived as more frustrating and/or mentally demanding. Interestingly, the per-

Table 3: General summary of the user study, including arithmetic means and standard deviations.

Group	<i>n</i>	Fails	<i>t</i>		<i>R</i>		<i>BA</i>		<i>FDR</i>		<i>FOR</i>		<i>MD</i>		<i>TD</i>		<i>PF</i>		<i>FL</i>	
			M	SD	M	SD	M	SD	M	SD	M	SD	M	SD	M	SD	M	SD	M	SD
Initial	-	-	-	-	-	-	.872	-	.500	-	.000	-	-	-	-	-	-	-	-	-
Total	102	3	22	13	70.3	19.2	.912	.079	.216	.147	.029	.035	.3	1.6	-.6	1.5	1.1	1.6	-.1	1.7
Neutral	34	0	26	15	67.5	12.4	.919	.058	.210	.107	.026	.025	.9	1.4	-.8	1.5	.9	1.5	.5	1.7
Accept	34	1	19	10	66.1	19.5	.923	.053	.208	.093	.022	.026	.2	1.5	-.8	1.5	1.2	1.6	-.2	1.8
Reject	34	2	21	13	77.2	22.8	.894	.112	.229	.214	.039	.047	-.2	1.8	-.4	1.6	1.1	1.7	-.6	1.6

Note: *M* for arithmetic mean and *SD* for standard deviation. *n* for the participant count and *Fails* for executions in rejecting H_0 . *BA*, *FDR* and *FOR* for measuring the Access Review Problem (ARP). *t* for the time-consumption of the AR and *R* for the amount of rejected UPAs. *MD* (Mental Demand), *TD* (Temporal Demand), *PF* (Performance) and *FL* (Frustration Level) for the NASA TLX.

ceived temporal demand *TD* did not correlate with *t*, possibly due to a missing baseline of a “normal” review duration. The result data showed a strong positive correlation between the perceived performance *PF* and actual performance *BA* (total: .607, neutral: .635, accept: .605, reject: .639), and a negative one between *PF* and the error rates *FDR* (total: -.336, neutral: -.516, reject: -.422; accept: not significant) and *FOR* (total: -.541, neutral: -.568, accept: -.480, reject: -.603). Therefore, the deciders had a realistic estimation of their performance. The result data also showed significant negative correlations between *FL* / *MD* and *BA* as well as positive ones between *FL* / *MD* and the error rates *FDR* / *FOR*, each for some groups. However, the causality remains unclear if deciders who find the task more difficult experience more stress, more stressed deciders deliver poor results, or both. Figure 9 (Appendix) shows the Spearman correlation and the local regressions.

► **Key takeaways of our user study:** (i) Almost all deciders improved the ARP. (ii) The required time differed substantially but was unrelated to quality (*BA*). (iii) *N06: Choice Defaults* led to reduced time effort and stress perception. (iv) A default reject led to more rejects. (v) A simple *N06: Choice Defaults* did not affect quality (*BA*) significantly but influenced the number of rejects. In detail, however, some increase in false rejects is tolerable as false accepts legitimate excessive authorizations leading to a false sense of security. (vi) Deciders’ self-assessed performance correlates significantly with *BA*, indicating the deciders’ realistic self-assessment.

7 Discussion

7.1 Acceptance Bias

Participants of the user study tend to accept existing authorizations. Existing research already documents and analyzes over-granting in real-world scenarios [18, 47, 54]. However, such scenarios involve strongly imbalanced data (see expert interviews: $1 - 22.8\% = 77.2\%$ of authorizations are estimated to be correct), social implications (a revoke acts against the interests of a real person), and unequal visibility of the two error types. Erroneous revokes are detected quickly, and the decider alone is responsible, while erroneous accepts are not

immediately visible and all previous approvers share the responsibility for also not resolving the error. With an initial *SEN* of 50% and no personal repercussions, the study had none of these biases and made no implication that acceptance is favorable to revocation. Still, deciders accept authorizations too often, with an average accept rate of 57.8% in the neutral group (see Section 6.1). While the study data does not explain this behavior, a possible explanation might be that the status-quo bias discourages deciders from revoking [42]: Following a real-world scenario, the study description states that participants need to review *existing* authorizations, which would be revoked upon rejection. The existence of a general status quo bias could also explain the relatively weak effect of the default accept bias on the accept rate: Study participants with default accept or reject nudge configuration needed to change an existing preselection to make an active decision and are thus also confronted with a status quo bias. If a status quo bias is already the reason for over-accepting in the neutral group, the effect of the default accept nudge would only repeat an already present bias. In contrast, the default reject nudge creates a new status quo that nudges the deciders in the opposite direction. The explanation seems plausible based on the study results, as the accept rate of the default accept group is closer to the neutral group (58.7%), and the accept rate of the default reject group is closer to the actual 50% (51.8%).

7.2 Implications for Access Review Challenges

► **Decider motivation affects quality (C4):** As described in Section 6.2, the user study participants had a reasonable estimation of their own performance. The user study design is fair, with a planned execution time of 20-30 minutes and no hurdles for *N01: Information Translation* or *N03: Information Visibility*. Still, some deciders submitted results with relatively low quality. The correlations between perceived stress (*FL*, *MD*) and quality (*BA*, *FDR*, *FOR*) may also indicate that decider motivation was an important factor. It must be assumed that poor decider motivation contributes stronger in real-world scenarios with larger scale and poorer information basis, indicating that nudges targeting decider motivation (*N09-N13*) may be a valuable contribution to AR quality.

► **Longer reviews are more demanding (C1, C4):** The user study results showed significant correlations between the review duration t and the perceived stress (FL , MD), underlining the importance of a reasonable scale. While the user study already confirms that *N06: Choice Defaults* considerably reduces review time, *N05: Range & Composition* also seems promising. Choice architects should take care not to overwhelm deciders with too many decisions. Distributing review responsibilities to many instead of a few decision-makers might be helpful. Considering *N10: Commitment Facilitation* or splitting reviews into multiple suitable sub-reviews carried out at different times or limiting them to unreviewed or changed authorizations could also improve quality.

► **N06: Choice Defaults effectivity does not seem to depend on decision difficulty (C2, C4, C5):** We tried to assess whether the impact of *N06: Choice Defaults* depends on the difficulty of a decision. For this purpose, we grouped the user study decisions by the 160 UPAs and their respective study group (neutral, default accept, default reject), resulting in $3 * 160$ groups of 34 review decisions. We then calculated the error rate and standard deviation for the decisions in the neutral group as indicators of the decision difficulty or uncertainty of UPA. To measure the effect of the default accept nudge for any UPA, we subtract the number of accepts in the neutral group from the amount of accepts in the default accept group. The resulting difference is the amount of *additional* accepts achieved by the nudge. The default reject effectivity was calculated as equivalent to the difference of rejects in the neutral and default reject groups. A Spearman correlation test with a $\alpha = .05$ significance level showed no significant correlation between the indicators for a decision's difficulty and the amount of additional accepts or rejects. The lack of correlation indicates that the effectivity of *N06: Choice Defaults* does not directly depend on the difficulty of a decision.

► **Spammers are an error source (C4):** Unlike the user study but in reality, a ground truth of detecting low-quality AR results is not available. Hence, it is helpful to identify "spammers" (deciders actually not trying to achieve an ARP improvement). The user study results suggest two possible ways to determine low-quality AR results: (i) While the review duration t did not correlate significantly with the quality metrics, we found that for the $n = 6$ deciders only taking $t = 6$ minutes or less, the mean BA ($M = 77.1\%$, $SD = 22.1\%$) drops a considerably $\Delta = -14.1\%$ comparing to BA of all participants ($M = 91.2\%$, $SD = 7.9\%$). (ii) Two spammers acted obviously ignorant by blindly accepting or rejecting all authorizations. In real-world scenarios, it might be helpful to use thresholds that, when undercut, classify the review as spam. We do not propose to dismiss such results categorically: it could be correct to accept all authorizations, or a decider could be quick. However, such deciders could be explicitly addressed to improve their result quality, e.g., by applying a custom nudge (like *N13: Empathy Instigation*) or requesting another person to check their decisions.

Table 4: Virtual best and worst advice.

Group	n	R	BA	FDR	FOR
Initial	-	-	.872	.500	.000
Total	102	70.3	.912	.216	.029
Virtual Best Advice	34	71.2	.931	.178	.023
Virtual Worst Advice	34	72.1	.885	.238	.043

Note: n for the participant count and R for the mean of rejected UPAs. BA , FDR and FOR are means for measuring the Access Review Problem (ARP).

► **Deciders have the last say (C4):** We re-grouped the user study decisions to simulate reviews with only correct and only incorrect *N06: Choice Defaults* (compare smart defaults [4, 5]). In reality, every decider had to make 160 decisions, of which 80 were TP (should be accepted) and 80 were FP (should be removed). This means that the default accept group had a correct preselection for exactly 80 authorizations, whereas the default reject group had a correct preselection for the other 80 ones. By virtually re-grouping these decisions, we create two sets of $34 * 160$ decisions each, for which one contains only correct default preselections and the other contains only incorrect ones. We then calculated the quality metrics BA , FDR , and FOR for both groups. Unsurprisingly, the virtual best advice group scored a higher overall quality than each of the three real study groups with $BA = 93.1\%$, and the lowest error rates with $FDR = 17.8\%$ and $FOR = 2.3\%$. The virtual worst advice group scored worse than all real groups with $BA = 88.5\%$, $FDR = 23.8\%$, and $FOR = 4.3\%$. However, the virtual best advice group's results are closer to those of all real groups than a perfect result, for which BA would be 100% and both error rates would be 0%. Similarly, the virtual worst advice group did not perform terribly but, in fact, still achieved a mean improvement in the ARP. Results for both groups show that users are affected by the *N06: Choice Defaults* and that the quality of the applied nudge affects the quality of the AR result. However, deciders have the last say and may choose not to follow a default, attenuating the worst assumptions of some interviewed experts. Table 4 summarizes the figures for both virtual groups.

7.3 Two Undesired Responsibility Shifts

Real-world access reviews (without nudge support) assume reflective decision-makers in transparent environments, leading to two assumptions: reflection and transparency [11]. However, the expert interviews and the user study discard both assumptions. For the reflection assumption, experts report several instances of human errors (C4), and the user study shows that deciders are affected by *N06: Choice Defaults*. Additionally, the deciders make errors despite having all the necessary data (even for the best advice in Table 4). For the transparency assumption, experts report the troublesome endeavor to present the information needed (*N01-N03*, *N09*) as too many or too few details lead to an unclear big picture.

Hansen and Jespersen [19] evaluate ethical considerations for nudge applications by the nudge's transparency and the decider's reflective or automatic mode of thinking. As mentioned earlier, access reviews should strive for transparency and reflective decisions. Access reviews in the real world and those with nudges can fail one of these: the real-world access reviews can lack transparency, and the nudged ones can lack reflective choices. On the one hand, real-world access reviews force reflective decisions as overwhelmed deciders actively need to choose, leading to a lack of transparency and constructing an unpleasant ethical situation. While reflective choices make the deciders fully responsible for their actions, the sheer scale (C1) and frequency (C3) put so many decisions on the table that the actual big picture for the access review becomes non-transparent. Therefore, the deciders have to bear the responsibility for a volume of decisions above their capabilities as human decision-makers, raising ethical concerns. On the other hand, the access reviews with the *N06: Choice Defaults* stay more transparent but allow for less reflective decisions, leading to a responsibility split. As soon as scale (C1) and frequency (C3) make the deciders give up on reflective choices, the choice architect shares responsibility for the decision-makers adopting its defaults.

In summary, neither burdening the deciders with the responsibility of choices they do not comprehend nor splitting the responsibility between the choice architect and the deciders are desired responsibility shifts for access reviews.

7.4 Design Implications for Usability

Following Hansen and Jespersen [19], design implications for future access reviews (with digital nudges) involve facilitating meaningful decisions based on transparency and reflective choices. When applied properly, digital nudges empower deciders to make confident and meaningful decisions with transparent and honest guidance [19]. Most importantly, this implies perceiving access review deciders and their decisions not as hyper-rational but as human, including their strengths and flaws [21, 40]. In the following, we derive three implications for usability based on our results.

► **Partition meaningfully:** Several experts find *N05: Range & Composition* relevant as it allows for meaningful partitions of access review decisions. Partitions effectively mitigate the deciders' scale perception and give a context for grouped decisions. Additionally, this allows abstract decisions for the whole partition. For example, deciding to revoke all authorizations of a person can be one meaningful decision instead of rejecting each of its authorizations one by one. Our experts name meaningful ways to partition decisions within access reviews, e.g., people leaving an organization, specific applications, critical authorizations, known past changes, organization-specific attributes, or processes. Ways to determine these partitions can range from choice architects' or deciders' experience to AI-based clustering.

► **Apply partition-specific digital nudges:** Digital nudges can be applied individually and combined for each partition. Based on the expert interviews, various digital nudges are suitable. For example, *N06: Choice Defaults* can preselect accepting security-uncritical authorizations (e.g., utility software) or rejecting security-critical ones (e.g., server access). Additionally, security-critical authorizations can be highlighted with a warning by *N02: Information Salience*. Thus, digital nudges can improve each partition's usability to guide access review deciders, also considering individual organizational contexts.

► **Query performance perception:** In the user study results, we find a strong correlation in all groups for the objective quality metric *BA* and the deciders' performance self-assessment *PF*. It shows that our user study participants had a reasonable perception of their performance. In contrast, a real-world access review cannot determine *BA* easily, as the underlying ground truth is unknown. This implies querying the deciders' performance self-assessment (*PF*) can be a valid and easy-to-implement estimator for the access review's quality (*BA*).

In summary, transparent digital nudges can guide human decision-makers to make meaningful, confident, and reflective choices. While the positive and negative effects of nudging require careful consideration, their anticipated effects are useful and promising tools for access review designs.

8 Conclusion

In this paper, we investigated digital nudges for access reviews. We formalized the access review problem. Subsequently, we interviewed highly qualified IAM experts to map the expected effects of digital nudges on access review challenges. Furthermore, we conducted a user study with *N06: Choice Defaults*. We found its influence on deciders' behavior in revoking authorizations. Additionally, we achieve time savings (up to 24.3%) and lower frustration. A simple *N06: Choice Defaults* did not significantly influence the overall quality, but it can shift the decisions to more revokes. While these revokes cause some false rejects, false accepts would be worse as they create a false sense of security by legitimating excessive authorizations. For future work, we invite researchers to study the ARP, to investigate other digital nudges of Table 2 or their combinations, or to replicate this study with a larger sample size or smart defaults [4, 5]. In sum, digital nudges are a promising tool to improve access reviews but need careful application.

Availability

For transparency and future research, we make the case study, all collected data, and the analysis of the user study open-source (<https://github.com/AccessReview/Availability>). In detail, we publish the instructions and data set of the case study, participants' results ($n = 102$), their choices ($n = 16,320$), and the *R* code to replicate our statistical evaluations.

Acknowledgments

The German Federal Ministry of Education and Research supported the research leading to these results as part of the DEVISE project (<https://devise.ur.de>).

This work would not have been possible without the help of our 10 interviewed experts and 102 user study participants. The experts invested a total of 10 hours and the participants a total of 42 hours to support our endeavor. Thank you!

References

- [1] William C. Adams. *Conducting Semi-Structured Interviews*, chapter 19, pages 492–505. John Wiley & Sons, Ltd, 2015.
- [2] Marvin Auf der Landwehr, Maik Trott, and Christoph von Viebahn. Consumers choice? fostering sustainability in grocery deliveries through digital nudging. In *Twenty-Ninth European Conference on Information Systems (ECIS 2021)*, ECIS 2021, page 1–16. Association for Information Systems, 2021.
- [3] Theophilus Azungah. Qualitative research: deductive and inductive approaches to data analysis. *Qualitative Research Journal*, 18(4):383–400, Jan 2018.
- [4] Paritosh Bahirat, Yangyang He, Abhilash Menon, and Bart Knijnenburg. A data-driven approach to developing iot privacy-setting interfaces. In *Proceedings of the 23rd International Conference on Intelligent User Interfaces*, IUI '18, page 165–176, New York, NY, USA, 2018. Association for Computing Machinery.
- [5] Paritosh Bahirat, Martijn Willemsen, Yangyang He, Qizhang Sun, and Bart Knijnenburg. Overlooking context: How do defaults and framing reduce deliberation in smart home privacy decision-making? In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI '21, New York, NY, USA, 2021. Association for Computing Machinery.
- [6] Basel Committee on Banking Supervision. Basel III: A global regulatory framework for more resilient banks and banking systems, June 2011.
- [7] Thomas Baumer, Mathis Müller, and Günther Pernul. System for cross-domain identity management (scim): Survey and enhancement with rbac. *IEEE Access*, 11:86872–86894, 2023.
- [8] Michelle Berger, Elias Greinacher, and Linda Wolf. Digital nudging to promote energy conservation behavior - framing and default rule in a smart home app. In *Thirtieth European Conference on Information Systems (ECIS 2022)*, ECIS 2022, page 1–16. Association for Information Systems, 2022.
- [9] Kristoffer Bergram, Marija Djokovic, Valéry Bezençon, and Adrian Holzer. The digital landscape of nudging: A systematic literature review of empirical research on digital nudges. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, CHI '22, New York, NY, USA, 2022. Association for Computing Machinery.
- [10] Kelly Caine. Local standards for sample size at chi. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, CHI '16, page 981–992, New York, NY, USA, 2016. Association for Computing Machinery.
- [11] Ana Caraban, Evangelos Karapanos, Daniel Gonçalves, and Pedro Campos. 23 ways to nudge: A review of technology-mediated nudging in human-computer interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI '19, page 1–15, New York, NY, USA, 2019. Association for Computing Machinery.
- [12] Federal Financial Supervisory Authority (BaFin). Rundschreiben 05/2023 (BA) - Mindestanforderungen an das Risikomanagement - MaRisk, June 2023.
- [13] Sandro Franzoi and Jan vom Brocke. Sustainability by default? nudging carbon offsetting behavior in e-commerce. In *Thirtieth European Conference on Information Systems (ECIS 2022)*, ECIS 2022, page 1–15. Association for Information Systems, 2022.
- [14] Ludwig Fuchs and Günther Pernul. HyDRo – hybrid development of roles. In *Information Systems Security*, pages 287–302. Springer Berlin Heidelberg, 2008.
- [15] Ludwig Fuchs, Günther Pernul, and Ravi Sandhu. Roles in information security – a survey and classification of the research area. *Computers & Security*, 30(8):748–769, 2011.
- [16] Cristina Gena, Pierluigi Grillo, Antonio Lieto, Claudio Mattutino, and Fabiana Vernero. When personalization is not an option: An in-the-wild study on persuasive news recommendation. *Information*, 10(10), 2019.
- [17] Abdul Muqet Ghaffar and Thomas Widjaja. Framing as an app-design measure to nudge users toward infection disclosure in contact-tracing applications. In *Thirty-first European Conference on Information Systems (ECIS 2023)*, ECIS 2023, page 1–16. Association for Information Systems, 2023.
- [18] Sebastian Groll, Sascha Kern, Ludwig Fuchs, and Günther Pernul. Monitoring access reviews by crowd labelling. In Simone Fischer-Hübner, Costas Lambroudikis, Gabriele Kotsis, A. Min Tjoa, and Ismail

- Khalil, editors, *Trust, Privacy and Security in Digital Business*, pages 3–17, Cham, 2021. Springer International Publishing.
- [19] Pelle Guldborg Hansen and Andreas Maaløe Jespersen. Nudge and the manipulation of choice: A framework for the responsible use of the nudge approach to behaviour change in public policy. *European Journal of Risk Regulation*, 4(1):3–28, 2013.
- [20] Sandra G. Hart. Nasa-task load index (nasa-tlx); 20 years later. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 50(9):904–908, 2006.
- [21] Jonas Hielscher, Uta Menges, Simon Parkin, Annette Kluge, and M. Angela Sasse. “Employees who Don’t accept the time security takes are not aware Enough”: The CISO view of Human-Centred security. In *32nd USENIX Security Symposium (USENIX Security 23)*, pages 2311–2328, Anaheim, CA, August 2023. USENIX Association.
- [22] Linda Hill. How automated access verification can help organizations demonstrate HIPAA compliance: A case study. *J Healthc Inf Manag*, 20(2):116–122, 2006.
- [23] Vincent C. Hu, David Ferraiolo, Rick Kuhn, Arthur R. Friedman, Alan J. Lang, Margaret M. Cogdell, Adam Schnitzer, Kenneth Sandlin, Robert Miller, Karen Scarfone, et al. Guide to attribute based access control (abac) definition and considerations (draft). Technical report, National Institute of Standards and Technology, 2014.
- [24] Dennis Hummel and Alexander Maedche. How effective is nudging? a quantitative review on the effect sizes and limits of empirical nudging studies. *Journal of Behavioral and Experimental Economics*, 80:47–58, 2019.
- [25] Matthias Hummer, Sebastian Groll, Michael Kunz, Ludwig Fuchs, and Günther Pernul. Measuring identity and access management performance - an expert survey on possible performance indicators. In *Proceedings of the 4th International Conference on Information Systems Security and Privacy*, pages 233–240. SCITEPRESS - Science and Technology Publications, 2018.
- [26] Pooya Jaferian, Hootan Rashtian, and Konstantin Beznosov. To authorize or not authorize: Helping users review access policies in organizations. In *Proceedings of the Tenth USENIX Conference on Usable Privacy and Security*, SOUPS ’14, page 301–320, USA, 2014. USENIX Association.
- [27] Mathias Jesse and Dietmar Jannach. Digital nudging with recommender systems: Survey and future directions. *Computers in Human Behavior Reports*, 3:100052, 2021.
- [28] Eric J. Johnson, Suzanne B. Shu, Benedict G. C. Dellaert, Craig Fox, Daniel G. Goldstein, Gerald Häubl, Richard P. Larrick, John W. Payne, Ellen Peters, David Schkade, Brian Wansink, and Elke U. Weber. Beyond nudges: Tools of a choice architecture. *Marketing Letters*, 23(2):487–504, Jun 2012.
- [29] Shelia M. Kennison, Ian T. Jones, Victoria H. Spooner, and D. Eric Chan-Tin. Who creates strong passwords when nudging fails. *Computers in Human Behavior Reports*, 4:100132, 2021.
- [30] Sascha Kern, Thomas Baumer, Ludwig Fuchs, and Günther Pernul. Maintain high-quality access control policies: An academic and practice-driven approach. In Vijayalakshmi Atluri and Anna Lisa Ferrara, editors, *Data and Applications Security and Privacy XXXVII*, pages 223–242, Cham, 2023. Springer Nature Switzerland.
- [31] Sascha Kern, Thomas Baumer, Sebastian Groll, Ludwig Fuchs, and Günther Pernul. Optimization of access control policies. *Journal of Information Security and Applications*, 70:103301, 2022.
- [32] Stefan Meier, Ludwig Fuchs, and Günther Pernul. Managing the access grid - a process view to minimize insider misuse risks. In *11th International Conference on Wirtschaftsinformatik (WI2013)*, pages 1051–1065, 2013.
- [33] Christian Meske and Tobias Potthoff. The dinu-model - a process model for the design of nudges. In *European Conference on Information Systems*, pages 2587–2597, 06 2017.
- [34] Tobias Mirsch, Christiane Lehrer, and Reinhard Jung. Making digital nudging applicable: The digital nudge design method. In *International Conference on Information Systems*. AIS, 2018.
- [35] Robert Münscher, Max Vetter, and Thomas Scheuerle. A review and taxonomy of choice architecture techniques. *Journal of Behavioral Decision Making*, 29(5):511–524, August 2015.
- [36] OWASP Top 10 team. Owasp top10, 2021. Accessed: 11/15/23.
- [37] Simon Parkinson and Saad Khan. A survey on empirical security analysis of access-control systems: A real-world perspective. *ACM Comput. Surv.*, 55(6), dec 2022.
- [38] Alexander Puchta, Fabian Böhm, and Günther Pernul. Contributing to current challenges in identity and access management with visual analytics. In Simon N. Foley, editor, *Data and Applications Security and Privacy*

- XXXIII, pages 221–239, Cham, 2019. Springer International Publishing.
- [39] Alexander Puchta, Sebastian Groll, and Günther Pernul. Leveraging dynamic information for identity and access management: An extension of current enterprise iam architecture. In *Proceedings of the 7th International Conference on Information Systems Security and Privacy - ICISSP*, pages 611–618, Online Streaming, 2021. INSTICC, SciTePress.
- [40] Ita Ryan, Utz Roedig, and Klaas-Jan Stol. Unhelpful assumptions in software security research. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security, CCS '23*, page 3460–3474, New York, NY, USA, 2023. Association for Computing Machinery.
- [41] Pierangela Samarati and Sabrina Capitani de Vimercati. Access control: Policies, models, and mechanisms. In Riccardo Focardi and Roberto Gorrieri, editors, *Foundations of Security Analysis and Design*, pages 137–196, Berlin, Heidelberg, 2001. Springer Berlin Heidelberg.
- [42] William Samuelson and Richard Zeckhauser. Status quo bias in decision making. *Journal of risk and uncertainty*, 1:7–59, 1988.
- [43] Ravi S. Sandhu. Role-based access control. portions of this chapter have been published earlier in sandhu et al. (1996), sandhu (1996), sandhu and bhamidipati (1997), sandhu et al. (1997) and sandhu and feinstein (1994). In Marvin V. Zelkowitz, editor, *Advances in Computers*, volume 46 of *Advances in Computers*, pages 237–286. Elsevier, online, 1998.
- [44] Armando Schär and Katarina Stanoevska-Slabeva. Application of digital nudging in customer journeys - A systematic literature review. In *25th Americas Conference on Information Systems, AMCIS 2019, Cancún, Mexico, August 15-17, 2019*. Association for Information Systems, 2019.
- [45] Christoph Schneider, Markus Weinmann, and Jan vom Brocke. Digital nudging: Guiding online user choices through interface design. *Commun. ACM*, 61(7):67–73, jun 2018.
- [46] Daniel Servos and Sylvia L. Osborn. Current research and open problems in attribute-based access control. *ACM Comput. Surv.*, 49(4), jan 2017.
- [47] Bingyu Shen, Tianyi Shan, and Yuanyuan Zhou. Improving logging to reduce permission Over-Granting mistakes. In *32nd USENIX Security Symposium (USENIX Security 23)*, pages 409–426, Anaheim, CA, August 2023. USENIX Association.
- [48] Cass R. Sunstein. The council of psychological advisers. *Annual Review of Psychology*, 67(1):713–737, 2016. PMID: 26393867.
- [49] Barnabas Szaszi, Anna Palinkas, Bence Palfi, Aba Szollosi, and Balazs Aczel. A systematic scoping review of the choice architecture movement: Toward understanding when and why nudges work. *Journal of Behavioral Decision Making*, 31(3):355–366, 2018.
- [50] Richard H. Thaler and Cass R. Sunstein. *Nudge: Improving decisions about health, wealth, and happiness*. Nudge: Improving decisions about health, wealth, and happiness. Yale University Press, New Haven, CT, US, 2008.
- [51] United States Congress. Health Insurance Portability and Accountability Act of 1996, 1996.
- [52] United States Congress. Sarbanes-Oxley Act of 2002. Corporate responsibility, 2002.
- [53] Markus Weinmann, Christoph Schneider, and Jan vom Brocke. Digital nudging. *Business & Information Systems Engineering*, 58(6):433–436, Dec 2016.
- [54] Tianyin Xu, Han Min Naing, Le Lu, and Yuanyuan Zhou. How do system administrators resolve access-denied issues in the real world? In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pages 348–361, 2017.
- [55] Sarah Y. Zheng and Ingolf Becker. Checking, nudging or scoring? evaluating e-mail user security tools. In *Nineteenth Symposium on Usable Privacy and Security (SOUPS 2023)*, pages 57–76, Anaheim, CA, August 2023. USENIX Association.
- [56] Samira Zibaei, Dinah Rinoa Malapaya, Benjamin Mercier, Amirali Salehi-Abari, and Julie Thorpe. Do password managers nudge secure (random) passwords? In *Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022)*, pages 581–597, Boston, MA, August 2022. USENIX Association.
- [57] Samira Zibaei, Amirali Salehi-Abari, and Julie Thorpe. Dissecting nudges in password managers: Simple defaults are powerful. In *Nineteenth Symposium on Usable Privacy and Security (SOUPS 2023)*, pages 211–225, Anaheim, CA, August 2023. USENIX Association.

Appendix

A Expert Interviews

A.1 Interview Script

I. Intro Section

Interview partner

- *What is your job position at organization XY?*
- *What is your IAM experience (years, clients, access review projects, managed identities, etc.)?*

Access review and its problems

- *Estimate the ratio of excessive granted access.*
- *Name 2-3 major challenges for access reviews.*

II. Explanation Section

- *Explain to the participant the access review challenges of Jaferian et al. [26]. Connect them to the major challenges of access review the participant named before.*
- *Explain to the participant digital nudges in general.*

III. Workshop Section

Mapping digital nudges and access review challenges

For each nudge in Table 5

1. *Explain the nudge and give an example fitting for the interview participant's environment.*
2. *The participant then freely reflects on the digital nudge and their relationship on access review challenges.*
3. *Finally, the participant rates each access review challenge, anticipating a very positive (+2), positive (+1), neutral (0), negative (-1), or very negative (-2) effect.*

Table 5: Digital Nudges [27] presented to the experts for mapping them to access review challenges [26].

Nudges	C1	C2	C3	C4	C5
Decision Information					
N01: Information Translation					
N02: Information Saliency					
N03: Information Visibility					
N04: Information Phrasing					
Decision Structure					
N05: Range & Composition					
N06: Choice Defaults					
N07: Option Consequences					
N08: Option-related Effort					
Decision Assistance					
N09: Reminders					
N10: Commitment Facilitation					
Social Decision Appeal					
N11: Messenger Reputation					
N12: Social Reference Point					
N13: Empathy Instigation					

Wrap-up

- *Name your TOP 3 digital nudges benefiting access review challenges.*

A.2 Codebook

We apply deductive and inductive coding to the expert interviews. The feasibility of digital nudges (based on the collection of Jesse and Jannach [27]) for access reviews suffice as interview questions. The access review challenges of Jaferian et al. [26] suffice as deductive codes, which we applied a priori to the interviews. Therefore, we trained and asked the interview partners about these challenges and asked for a Likert scale-based rating (2 (best), 1, 0, -1, -2 (worst)). The experts answered with different arguments, for which we extracted inductive codes. The rating for digital nudge, challenge and the inductive codes are detailed in the codebook (Table 6).

B User Study

B.1 Data Set

For the user study, we used a crafted data set (160 UPAs). We can pinpoint which UPAs are correctly (TP) and incorrectly (FP) assigned. Figure 5 (using a grid representation based on [32]) depict the data set. A processable format is available at GitHub.

B.2 Ground Truth Document

Access Review Case Study

You work as a busy head of the marketing department in a large industry company with many concurrent projects to maximize the income for your company. Your time is limited, and you have marketing goals to fulfill.

The security teams reminded you via email that your company is legally required (compliance) to review the permission assignments for the employees in your department. You must follow the **principle of least privilege**: Employees must have permissions required for their job, but not more. If you decide to revoke an excessive permission for one of your employees, the employee will no longer be able to access the associated resources by tomorrow.

While the security team points out that any excessive permission poses a security threat, you are aware that missing ones might prevent your employees from working until they re-obtain it via a time-consuming help-desk or self-service request.

The marketing department consists of three teams:

I. Graphic design team

- Create and edit images for the company's media and advertisement presence. This includes banners, logos, websites, or campaign designs that are used in advertisements or social media posts.
- Require a Photoshop license to work.

II. Social media team

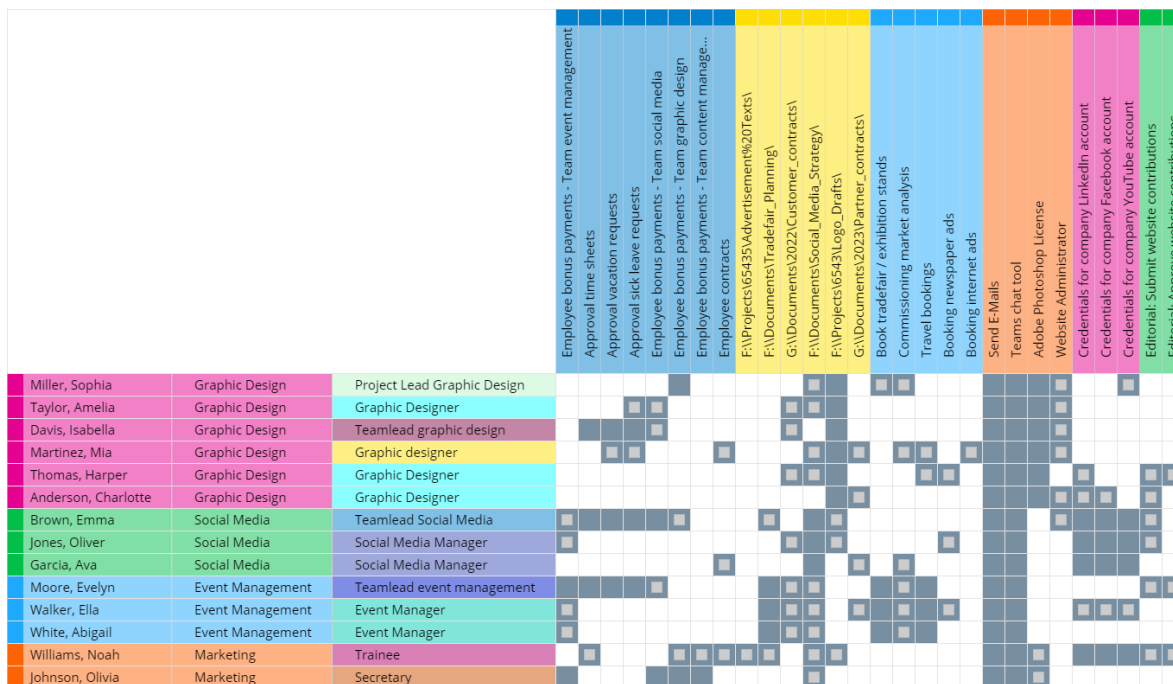


Figure 5: Grid visualization [32] of the user study data set. Blue cells resemble TP, gray ones FP, white ones TN, and FN were not present in the data set.

- Manage the company’s social media accounts.
- Need to communicate with potential customers, candidates for recruiting, and partners online.

III. Event management team

- Organize trade fairs and partner events across West and Central Europe.
- Book trade fair stands.
- High self-organization; often need to attend remote events without long preparation.

IV. Department hierarchies

- Every team is led by a team lead who overlooks the employee’s attendance and work results.
- Team leads have an annual budget for bonus payments, which they can distribute among their team members based on last year’s performance. The secretary reads the specified bonus payments defined by the team leads from the HR system and arranges for the salary to be posted.

- The department’s trainee used to intern in the graphic design team. Now, he is working in the social media team.

V. Misc

- Everybody communicates with MS Teams and Outlook.
- You can sort the columns.

B.3 Screenshots Access Review

We used three configurations of the access reviews with the same data basis. Figure 6, the neutral default, has two white buttons without a preselection. Figure 7 displays the default accept with a preselected *Approve*. Figure 8 shows the default reject with a preselected *Remove*.

B.4 Statistical Analysis

Figure 9 depicts a pair plot for each metric separated for their group. Green shows the default accept group, red the default reject, and blue the neutral one. The upper right part depicts Spearman correlations. The stars indicate the significance levels as "****": $p < .001$; "***": $p < .01$; "**": $p < .05$, and "." $p < .1$. The lower left depicts local regressions. Finally, the diagonal, the first row and column show metric distributions.

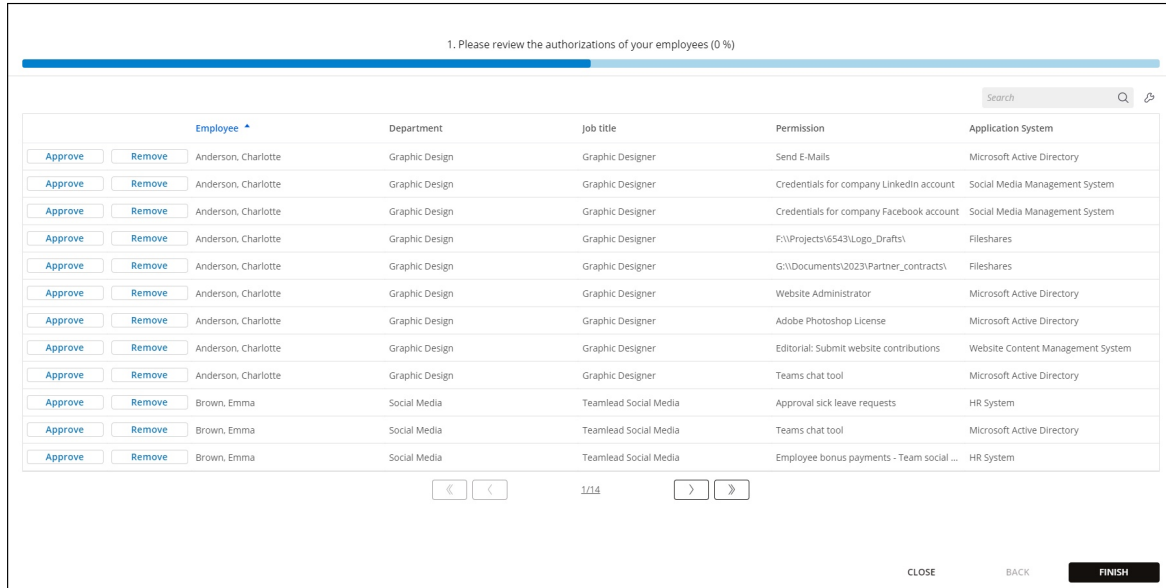


Figure 6: Screenshot of the neutral group for the user study.

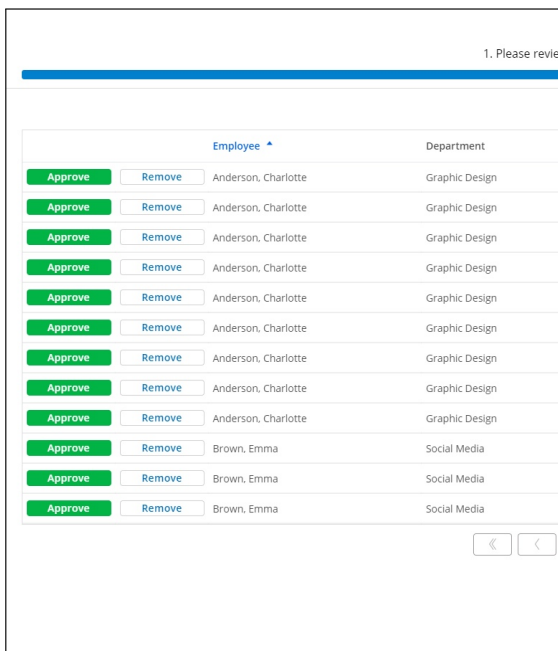


Figure 7: Screenshot for the accept group.

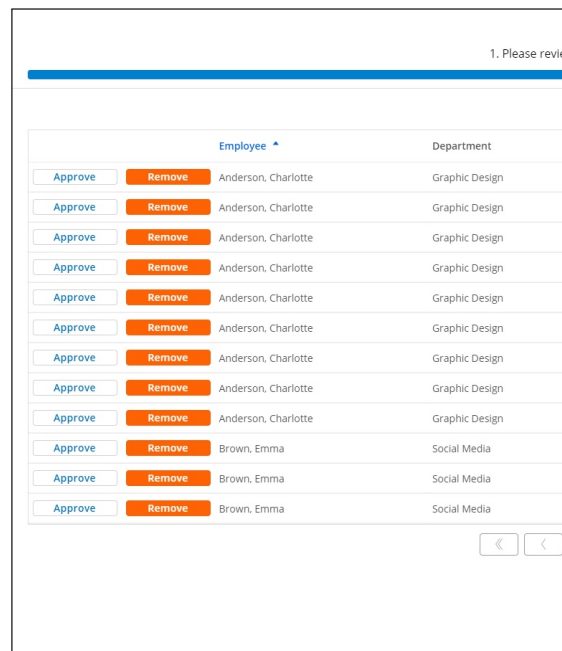


Figure 8: Screenshot for the reject group.

Table 6: Codebook for expert interviews.

N	C	Likert	Inductive Codes
N01	C1	1	Understandability (E02, E03, E04, E09); No effect (E03, E06, E08, E10); Feel-Good (E02, E03); Uniqueness (E04, E09); Structure (E09)
N01	C2	2	Understandability (E02, E05, E06, E07, E08, E09, E10); Mental Load (E03, E05, E07, E09); Acceptance (E05, E07, E10); Wording (E05, E06, E07)
N01	C3	1	Recognition (E01, E04, E05, E06, E09); Learning (E04, E05, E09); Feel-Good (E05, E06)
N01	C4	2	Understandability (E02, E05, E07, E08, E09, E10)
N01	C5	0	Understandability (E05, E06); Recognition (E04)
N02	C1	1	Focus (E01, E02, E04, E05, E09); No effect (E06, E08, E10)
N02	C2	0	Focus (E06, E09, E10); No effect (E03, E07, E09)
N02	C3	1	Economic Efficiency (E01, E02, E04, E07, E09); Focus (E01, E02); Acceptance (E09)
N02	C4	1	Focus (E03, E05, E06, E07, E08, E09, E10); Algorithm-Quality (E03, E06, E09); Backlash (E03, E06, E09)
N02	C5	2	Focus (E03, E04, E05, E07, E10); Algorithm-Quality (E09)
N03	C1	1	More relevancy (E03, E04, E05, E07, E09); Less confusion (E03, E04, E05, E09); No reduction of decisions (E06, E08)
N03	C2	2	Showing more data (E01, E05, E08, E09, E10); Relevancy (E03, E04, E06)
N03	C3	0	Run-time (E05); Recognition (E07)
N03	C4	1	Mistake mitigation (E05, E07, E09, E10); Focus (E06, E07)
N03	C5	2	Showing more data (E01, E07, E09, E10); Need to know (E07, E10)
N04	C1	0	Insecurities of decision-maker (E09); Sense of responsibility (E07)
N04	C2	-1	Context-Awareness (E05, E07, E08, E09, E10); Bias (E04, E05, E09), Base direction (E05, E09)
N04	C3	0	Acceptance (E07)
N04	C4	1	Acceptance (E06, E08, E09, E10); Focus (E06, E09, E10); Pressure (E02)
N04	C5	0	Focus (E06, E07, E10)
N05	C1	2	Similarities (E01, E03, E04, E05, E07, E08, E09); Overhead (E08, E10)
N05	C2	1	Focus (E04, E08, E09, E10); Audience (E08, E09)
N05	C3	1	Economic Efficiency (E03, E05, E06); More Tasks (E09, E10)
N05	C4	2	Focus (E01, E03, E05, E06, E07, E09, E10); Similarities (E01, E05, E06, E07); Smaller Batches (E09, E10)
N05	C5	2	Exceptional Case Detection and View (E02, E03, E04, E07, E09)
N06	C1	2	Less work (E01, E02, E04, E05, E06, E09); No reduction of decisions (E07, E10)
N06	C2	-2	Recommendation Fallacy (E02, E04, E05, E06, E07, E09, E10); Recommendation Support (E06, E09)
N06	C3	2	Less work (E01, E02, E04, E05, E06, E08, E09)
N06	C4	-2	Less diligence/Focus (E01, E02, E04, E05, E06, E07, E09, E10); Recommendation Fallacy (E02, E04, E05, E06, E07, E09, E10)
N06	C5	0	Not in Focus (E05, E07, E10); Default handling (E06, E09); Special treatment (E04)
N06	Misc		Not Compliant (E01, E03, E07, E09); Needs good recommendation (E01, E03, E08, E09); Is it really a decision? (E03, E07, E09)
N07	C1	0	Speed (E01, E09)
N07	C2	-1	Recommendation Fallacy (E09)
N07	C3	1	Speed (E01, E04, E09); Gamification (E04, E05, E09); Feel-Good (E04, E07); Acclimatation (E09)
N07	C4	-1	Pressure (E01, E03, E07, E09, E10); Recommendation Fallacy (E06, E07, E09, E10); Less diligence (E01, E03, E07)
N07	C5	-1	Recommendation Fallacy (E07, E09); Fairness for disadvantaged individuals (E09)
N08	C1	-1 / 1	Ambivalence (E01, E03, E05, E06, E07, E08, E09, E10); Economic Efficiency (E02, E03, E05, E07, E08); Acceptance (E04, E07, E09, E10)
N08	C2	1 / -1	Ambivalence (E01, E03, E05, E06, E07, E08, E09, E10)
N08	C3	-1 / 1	Ambivalence (E01, E03, E05, E06, E07, E08, E09, E10); Economic Efficiency (E02, E03, E05, E07, E08); Acceptance (E04, E07, E09, E10)
N08	C4	1 / -1	Ambivalence (E01, E03, E05, E06, E07, E08, E09, E10); Economic Efficiency (E02, E03, E05, E07, E08); Acceptance (E04, E07, E09, E10)
N08	C5	1 / -1	Ambivalence (E01, E03, E05, E06, E07, E08, E09, E10)
N09	C1	0	No effect (E03, E07); More participation (E04, E10);
N09	C2	1	Instructions and Goals (E02, E03, E04, E05, E07, E08, E09, E10); Spam (E01, E02, E03, E07, E09, E10);
N09	C3	2	Spam (E01, E02, E03, E07, E09, E10); Attention (E04, E06, E07, E09, E10)
N09	C4	-1	Revisit (E03, E07, E08); Pressure (E05); Multi-Channel (E07, E09, E10); Audience (E03, E09)
N09	C5	0	Open Task (E01); No effect (E07)
N10	C1	1	Combination with N05 - Commitment for partitions (E02, E04, E06, E07, E08, E09, E10)
N10	C2	0	Autonomic planning and understanding (E04, E05, E08, E09)
N10	C3	1	Combination with N05 - Sub-Deadlines for partitions (E05, E07, E08, E09, E10); Comfort (E02, E06, E07, E08, E09, E10)
N10	C4	1	Focus (E04, E06, E07, P8, E10); Comfort (E02, E06, E07, E08, E09, E10)
N10	C5	0	Focus (E07, E10)
N11	C1	1	Endurance (E01, E02, E03, E04, E05, E07, E09, E10); Trust (E02, E03, E07, E08, E09)
N11	C2	2	Approachable IAM team (E01, E02, E03, E04, E05, E07, E08, E09, E10)
N11	C3	1	Endurance (E01, E02, E03, E04, E05, E07, E09, E10); Trust (E02, E03, E07, E08, E09)
N11	C4	2	Approachable IAM team (E01, E02, E03, E04, E05, E07, E08, E09, E10); Acceptance (E01, E02, E03, E04, E08, E10)
N11	C5	2	Approachable IAM team (E01, E02, E03, E04, E05, E07, E08, E09, E10)
N12	C1	0	Endurance (E02, E03, E06, E09); Backlash (E09)
N12	C2	2	Approachable Peer-Group (E02, E03, E04, E07, E08, E09, E10)
N12	C3	0	Endurance (E02, E03, E06, E09); Backlash (E09)
N12	C4	1	Acceptance (E02, E03, E06, E07, E09, E10); Peer-Pressure (E03, E10)
N12	C5	2	Approachable Peer-Group (E02, E03, E04, E07, E08, E09, E10)
N12	Misc		Similarity to N11 messenger reputation (E01, E05)
N13	C1	1	Feel-Good (E02, E04, E06, E07, E09)
N13	C2	1	Feedback on odd behavior (E02, E03, E04, E07, E09)
N13	C3	1	Feel-Good (E02, E04, E06, E07, E09)
N13	C4	1	Feel-Good (E02, E05, E07, E08); Focus (E02, E04, E07, E09)
N13	C5	0	Feel-Good (E05)

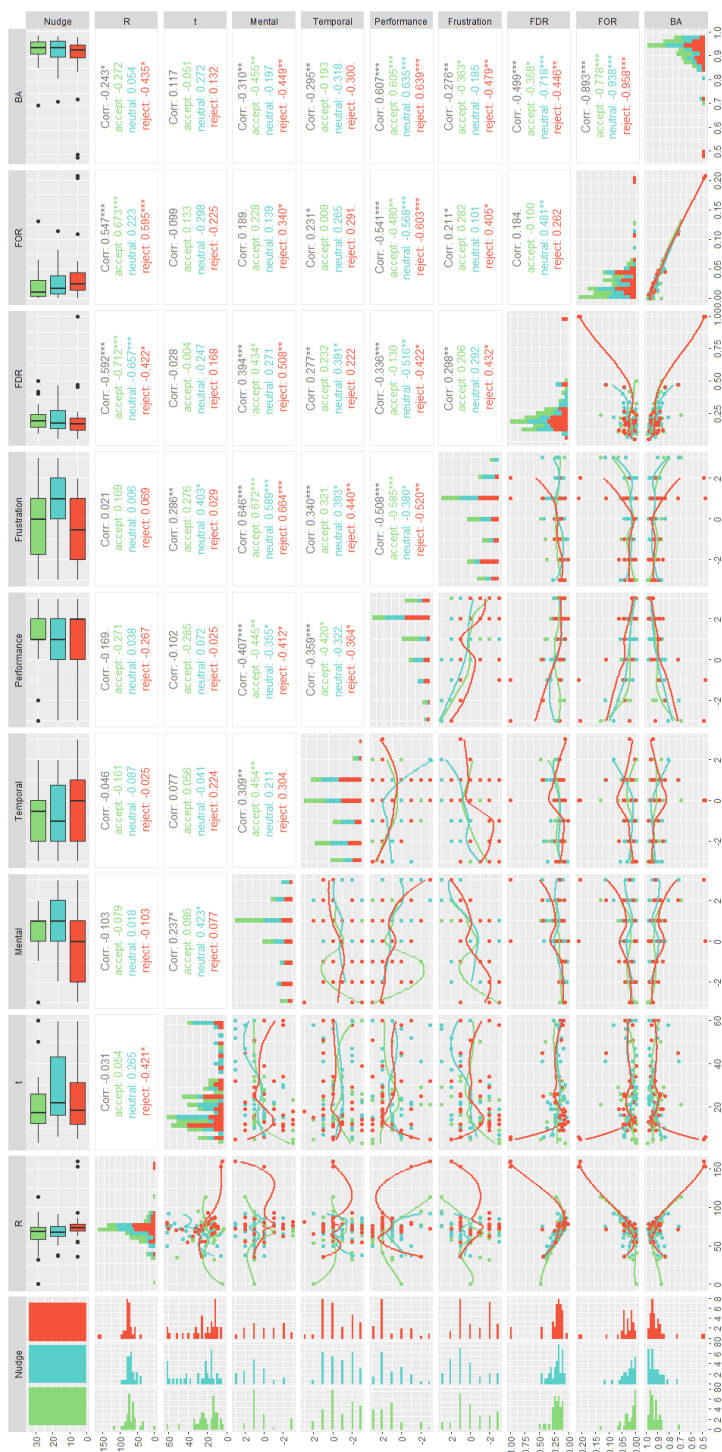


Figure 9: Pair plot of correlations (Spearman) and local regressions for the user study.