


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) [47,35] 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>).

an organization or its IT infrastructure, incorrect policy updates and a common practice of granting permissions too freely [47] cause ACPs to deteriorate and lose quality over time [45,23]. 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)* [37], 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) [39], Role-Based Access Control (RBAC) [13,38] and Attribute-Based Access Control (ABAC) [25] 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 [34,3,28]. 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 [29].

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 [11,44], policy mining algorithms evaluate existing permission assignments to generate new policies based on them [32,48]. Hybrid approaches try to combine the advantages of both types [17]. 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 [27]. 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 ABAC policies [20]. 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 [24]. 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 [4,30]. 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, which are listed in table 5. 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

Table 1. Considered literature.

Literature	P1	P2	P3	P4	P5
L1: Jaferian et al. [29]	X	X	X	X	
L2: Puchta et al. [36]	X	X		X	X
L3: Parkinson and Khan [35]	X	X			
L4: Servos and Osborn [40]	X	X		X	X
L5: Smetters and Good [41]	X	X			
L6: Fuchs et al. [15]	X	X			
L7: Hummer et al. [27]	X	X			X
L8: Groll et al. [18]	X	X	X	X	
L9: Hill [22]		X		X	
L10: Benedetti and Mori [8]			X	X	
L11: Hu et al. [24]	X		X		
L12: Strembeck [42]		X			
L13: Xu et al. [47]	X				
L14: Xiang et al. [46]		X	X		
L15: Bauer et al. [5]	X	X			
L16: Kunz et al. [31]	X	X			X
L17: Kern et al. [30]	X	X			X

and issues in respect to ACPs. However, we are going to give a general classification of their 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 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 and an IAM engineer 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

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	Pharmaceutics	X	X	X	X	X
EI3: IAM officer, IAM engineer	Insurance	X	X	X	X	X
EI4: IAM officer, IAM engineer	Retail	X	X	X	X	X
EI5: CISO	Software & Consulting		X	X	X	
EI6: 2 IAM consultants	Software & Consulting	X	X	X	X	X

also not desirable due to the role explosion problem [14]. Moreover, all companies 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. outdating 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 struc-

tures 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 [18]. 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 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 [29,47]. 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" [16,26], 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 [31]. 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 [21] 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 [26]. 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 [30]. Compliance goals typically overlap

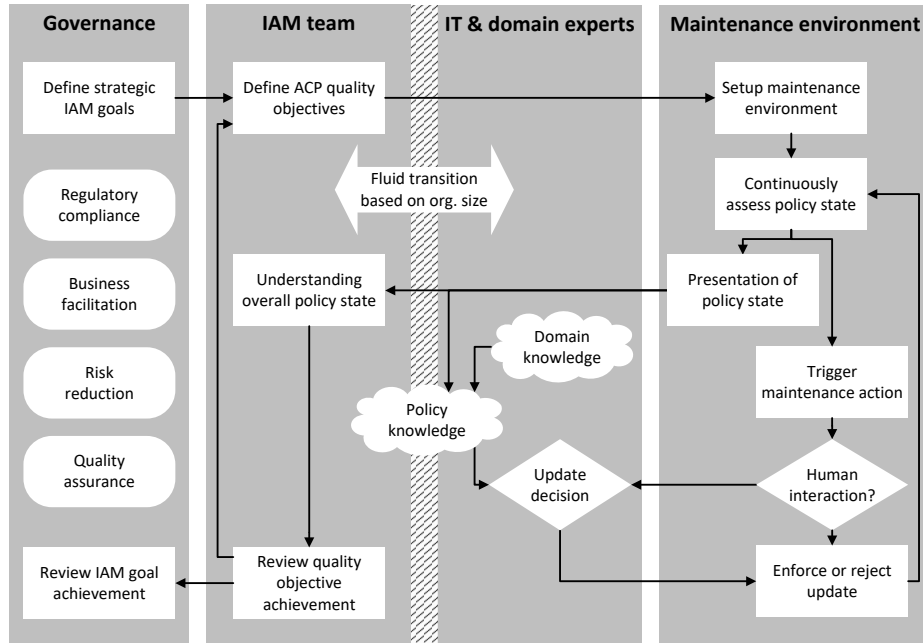


Fig. 1. Schematic overview of the proposed ACP Maintenance Framework.

with the aforementioned, and may also include adherence to the principle of least privilege, implementation of SoD policies, definition and adherence to formal responsibilities, and comprehensive change logs and reporting capabilities to verify compliance with these requirements. Based on the formulated strategic goals, the IAM team evaluates the available resources and defines ACP quality objectives that can be achieved in a given period (e.g., 6 months). These quality objectives serve as basis for the ACP maintenance process and must be checked and reviewed periodically. Since maintenance capabilities must be implemented and evolved over time, the initial quality objectives must be formulated at a low level. As the maintenance process matures, they can then be gradually increased until a satisfactory ACP quality level is reached and can be maintained over the long term. The IAM team reports to the governance domain using appropriate key performance indicators so that the governance domain can monitor the achievement of the formulated strategic goals.

4.2 Implementing the ACP maintenance environment

The ACP maintenance is supported by a maintenance environment. This can include, for example, data analysis and visualization tools, workflow tools, IAM-specific tools or in-house developments. At the beginning of the implementation, an understanding of the existing entitlement structures must be achieved. Since the IAM-relevant data is typically distributed in a heterogeneous form in a large

number of applications (e.g. target systems, IAM systems, supplementary data 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. [31]). 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 [43,12,10]. 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 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 struc-

ture, 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 [33]. 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 [19].

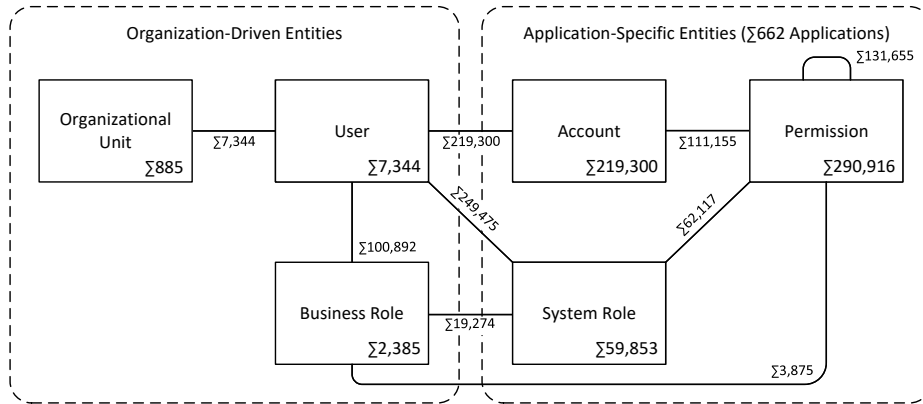


Fig. 2. Assumed data model of the case study with initial entity counts.

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 [31]. 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 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.

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., Vaniea, 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>

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. Fuchs, L., Pernul, G.: HyDRo – hybrid development of roles. In: Information Systems Security, pp. 287–302. Springer Berlin Heidelberg (2008). https://doi.org/10.1007/978-3-540-89862-7_24, https://doi.org/10.1007/978-3-540-89862-7_24
18. 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)
19. 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. pp. 161–172 (2013)
20. 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>
21. 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)
22. Hill, L.: How automated access verification can help organizations demonstrate HIPAA compliance: A case study. *J Healthc Inf Manag* **20**(2), 116–122 (2006)
23. 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)
24. 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)
25. 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>
26. 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>
27. 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>
28. 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
- 29. 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)
- 30. 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>
- 31. 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>
- 32. Mitra, B., Sural, S., Vaidya, J., Atluri, V.: A survey of role mining. *ACM Computing Surveys (CSUR)* **48**(4), 1–37 (2016)
- 33. 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)
- 34. 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
- 35. 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>
- 36. 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)
- 37. 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)
- 38. 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>
- 39. Sandhu, R.S., Samarati, P.: Access control: principle and practice. *IEEE communications magazine* **32**(9), 40–48 (1994)
- 40. 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>
- 41. 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>
- 42. Strembeck, M.: Scenario-driven role engineering. *IEEE Security & Privacy* **8**(1), 28–35 (Jan 2010). <https://doi.org/10.1109/MSP.2010.46>

43. 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)
44. 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)
45. Xia, H., Dawande, M., Mookerjee, V.: Role refinement in access control: Model and analysis. *INFORMS Journal on Computing* **26**(4), 866–884 (2014)
46. 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>
47. 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>
48. Xu, Z., Stoller, S.D.: Mining attribute-based access control policies. *IEEE Transactions on Dependable and Secure Computing* **12**(5), 533–545 (2014)

A Appendix

Table 5. Question catalogue for the semi-structured expert interviews

Interview Question
Q1: Which ACMs are used in your organization?
Q2: Who is responsible for maintaining ACPs in daily operations?
Q3: Who in your organization understands meaning & effect of selected ACPs?
Q4: Do you perform access reviews? If yes, what exactly is being reviewed?
Q5: What motivation and goals would speak for improving existing ACP structures?
Q6: Which efforts do you make to improve ACP structure?
Q7: (How) do you prioritize which ACPs to improve?
Q8: How regularly does structural improvement take place?
Q9: How clearly is this structured and documented?
Q10: Which components of structural improvement are currently automated?
Q11: What degree of automation would be desirable or realistic?
Q12: Are there examples of specific problems encountered during maintenance or optimization of existing ACP structures?
Q13: What are the major structural challenges in ACP maintenance?