A Trust and Reputation System for Examining Compliance with Access Control

Thomas Baumer
thomas.baumer@nexis-secure.com
Nexis GmbH
Regensburg, Germany

Johannes Grill
johannes.grill@informatik.uni-regensburg.de
Universität Regensburg
Regensburg, Germany

Jacob Adan
jacobadan562@gmail.com
Universität Regensburg
Regensburg, Germany

Günther Pernul
guenther.pernul@informatik.uni-regensburg.de
Universität Regensburg
Regensburg, Germany

ABSTRACT

Trust is crucial when a truster allows a trustee to carry out desired services. Regulatory authorities thus set requirements for organizations under their jurisdiction to ensure a basic trust level. Trusted auditors periodically verify the auditee's compliance with these requirements. However, the quality of the auditees' compliance and the auditors' verification performance often remain unclear and unavailable to the public. In this work, we examine the regulations of Identity and Access Management (IAM) and identify typical patterns. We enhance these patterns to include trust measurements for the auditee providing services and the auditors verifying compliance. We demonstrate the feasibility of this approach for an application utilizing decentralized blockchain technologies and discuss the implications, potential, and benefits of this architecture.

CCS CONCEPTS

• **Computer systems organization** → **Dependable and fault-tolerant systems and networks**; • **Information systems** → **Reputation systems**; • **Security and privacy** → **Trust frameworks**; **Access control**; • **Social and professional topics** → **Technology audits**; **Computing / technology policy**.

KEYWORDS

Trust, Identity and Access Management, Regulation, Trust and Reputation Systems, Blockchain

ACM Reference Format:
Thomas Baumer, Johannes Grill, Jacob Adan, and Günther Pernul. 2024. A Trust and Reputation System for Examining Compliance with Access Control. In *The 19th International Conference on Availability, Reliability and Security (ARES 2024), July 30-August 2, 2024, Vienna, Austria.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3664476.3670883

ARES 2024, July 30-August 2, 2024, Vienna, Austria

ACM ISBN 979-8-4007-1718-5/24/07
https://doi.org/10.1145/3664476.3670883

1 INTRODUCTION

Identity and Access Management (IAM) is a cornerstone for modern IT security as it protects digital identities and resources by proper access control. Regulative authorities acknowledge its relevance and demand proper controls to ensure IAM essentials. These regulations include Sarbanes-Oxley Act (SOX) [52], Basel III [3], or electronic IDentification, Authentication and trust Services (eIDAS) [13]. In most cases, an auditor periodically verifies that an auditee complies with relevant regulations. Penalties apply to non-compliant auditees for failure to comply with the regulation. For the public domain, a failing audit usually only becomes relevant if penalties lead to severe economic consequences like exclusion from certain markets.

Theoretically, this should ensure sufficient cybersecurity, and organizations report these compliance requirements as crucial motivators for cybersecurity [28]. However, we can observe failures of these regulations in scandals or corporate collapses. Prominent and recent examples with remarkable consequences (not only for cybersecurity) are present: (i) The Sarbanes-Oxley Act (SOX) results from accounting scandals like Enron and WorldCom. Enron and its auditor, Arthur Anderson LLP, collapsed in the following events [22]. This collapse depicts the downfall of one of the five largest auditor companies, reducing the *Big Five* to the now notable *Big Four* (EY, Deloitte, KPMG, PwC). (ii) More recently, the Wirecard scandal displays Germany's biggest accounting fraud in post-war history. Teichmann et al. conclude on Wirecard: *"the company's internal monitoring, third-party audits and state monitoring by BaFin [the regulative authority] were inadequate"* [50]. Meanwhile, EY, Wirecard's auditor, lost some of its reputation and market share while legal actions were still ongoing at the time of writing this paper. (iii) Deloitte was fined $1.1 million for altering timestamps on backdating audit documents for obscuring audit processes, showcasing integrity issues [27]. (iv) A few weeks ago, the auditor KPMG set the record for a $25 million fine because its employees cheated on exams by sharing questions and answers for mandatory professional certificates to conduct audits. Ironically, the queried topics included professional ethics [33]. (v) In Australia, the auditor PwC simultaneously consulted Australian enterprises and their government. PwC leaked the government's tax plans to enterprises, helping them avoid tax laws that PwC had drafted themselves [41]. In summary, while regulations motivate audited organizations to take action toward cybersecurity [28], scandals

across large audit companies raise the question: *besides the auditees, can we trust the auditors and can we measure the trust for both?*

We propose using a decentralized trust and reputation system for audits involving IAM. This system needs to be transparent and tamper-proof as the trust towards the auditees and auditors is limited in the first place, while the privacy needs of the involved parties need consideration. A decentralization using blockchain technologies further enhances security properties regarding integrity, availability, and avoidance of a single point of failure. Therefore, we use the following research question to facilitate our research:

RQ *How to transparently and tamper-proof express trust and reputation for auditees and auditors to act securely according to IAM regulations while respecting their privacy needs?*

To answer this research question, we utilize the following method. Our primary goal is to design a trust and reputation system for documenting trust in compliance with IAM regulations. Therefore, we first study relevant IAM regulations and then design a secure architecture. We evaluate this architecture's performance, typical attacks toward trust and reputation systems, and a use case. A final discussion wraps up the findings for future research. This method leads to the following contributions:

- We propose a transparent architecture for making the trustworthiness of auditees and auditors accessible.
- We measure a sufficient efficiency for major IAM regulations and make our prototype open-source available.[1]
- We argue resistance to acknowledged security concerns for trust and reputation systems and examine a use case.
- We point out options to calculate trust, applying the trust score in end-user-friendly badges, and that our trust and reputation system is not limited to IAM.

We outline the remaining work as follows. Section 2 introduces the background of this work, including terminology and related work. Section 3 details on the applied methodology. Subsequently, Section 4 designs the architecture of the trust and reputation system, evaluated in Section 5. Section 6 discusses further insights of this work. Finally, Section 7 concludes this paper.

2 BACKGROUND

In this Section, we cover the terminology used within this work for trust, identity and access management, and blockchain. Furthermore, we provide an overview of related work.

2.1 Terminology

➤ **Trust**: is a complex property with various definitions in the analogous and digital world. Trust is prevalent when a truster entrusts a trustee to provide a service, which may negatively affect the truster. An intuitive application of trust to the digital world is *computational trust*. This abstraction of trust allows for a more measurable notion of trust, easing decision-making. *Reputation*, similarly, is a widespread concept without a precise definition. It describes the perception a party has of another, for example, based on past transactions. Like trust, *computational reputation* makes reputation measurable for decision-making. *Trust and Reputation Systems* build upon these notions and are collaborative systems that enable or sanction their users based on their accumulated trust and reputation. These trust and reputation systems are essential for modern cybersecurity as they allow for investigating the reputation of a trustee before entrusting a service execution. A typical commercial example of these systems is www.ebay.com since its participants must trust their reputations for successful deals. [7, 26]

➤ **Identity and Access Management (IAM)**: is a cornerstone of modern cybersecurity as it includes processes, policies, and technologies for introducing and maintaining proper controls for authentication and authorization [16]. Therefore, it combines various research topics, including identity management [4, 9], access control [17, 29, 34, 43], and authentication [32, 38]. Regulative authorities also acknowledge the relevance of IAM and demand internal controls for proper authentication and authorization, which effectively translate to an implementation need for IAM systems. Therefore, the effective fulfillment of IAM requirements also intersects with IT management and corporate governance.

➤ **Blockchain**: is a decentralized database in which participants who do not trust each other can agree on the state of the database [48]. A blockchain comprises multiple nodes, each storing a replica of the database [48]. It facilitates transparency and data integrity by executing decentralized program logic on all nodes, so-called *Smart Contracts*, with the outcome being determined by majority consensus [11]. Each participant holds a cryptographic key pair. The private key is used to sign transactions on the blockchain. The public key verifies the authenticity of transactions [11]. Blockchain can further be categorized into different types [5]. In *public-permissionless* blockchains, anyone can establish an identity, engage in transactions, and access all transaction records. Conversely, in *private-permissioned* blockchains, only designated participants can create an identity, conduct transactions, and access transaction histories.

2.2 Related Work

Trust and, more precisely, trust and reputation systems are widespread research topics across cybersecurity disciplines. Thus, various surveys summarize work on trust and reputation systems [1, 6, 7, 19, 26, 54]. Recent advances in trust usually concentrate on mitigating trust, using trust by relying on a trusted third party, or managing trust. This work focuses on managing trust in socio-technical systems, which renders related work on trust for cryptography or hardware security out of the scope.

Recent application domains include models for assessing whether an Artificial Intelligence (AI) can be a trustee when it executes a service for a truster [25]. Kuang et al. propose a trust and reputation system using blockchain technologies to detect malicious nodes within vehicular networks [31]. Sheng et al. propose another blockchain improvement by assessing trust for environments with multiple blockchains [45]. As part of IAM realm, Tan et al. propose a multi-factor authentication system for distributed trust systems [49]. Nuss et al. also propose a design for IAM with blockchain technologies [37]. Another exciting application of trust and reputation systems is sharing communities for Cyber Threat Intelligence (CTI). Geras and Schreck name trust a crucial component for successful CTI sharing communities after interviewing several security and CTI experts [18]. Additionally, Sayeed et al. propose a decentralized

[1] https://github.com/TrustInRegulations/Code

system for CTI collection [44]. Trust and reputation are also relevant topics for privacy in cloud computing. Simou et al. [46] review privacy and trust issues in a cloud context, while Reijsbergen et al. [42] study transparent and privacy-aware data processing for cloud services.

In summary, while much recent work on trust and reputation systems exists, to the best of our knowledge, one focusing on IAM regulations with modern technologies and considering the trustworthiness of both auditees and auditors misses.

3 METHOD

This work aims to transparently and tamper-proof express trust for auditees and auditors based on the fulfillment of IAM regulations. Therefore, effective communication of this auditee trust is necessary while respecting the secrecy requirements of auditees and public accessibility for individuals using their services.

Therefore, we first research the relevant regulations for IAM. The goal is to extract common processes and overarching similarities. We search for IAM regulations by unstructured search queries with common search engines like Google Search[2] or Microsoft Bing[3]. Once we find a regulation, we evaluate whether it requires implementing controls to ensure proper authentication or authorization for including it in our result set. If that evaluation is not possible because of language barriers or limited availability in general, we exclude this specific regulation. We classify the IAM regulations based on the result set and extract common patterns. Section 4.1 details the resulting regulations, their classification, and patterns.

With these extracted classifications and patterns, we model a general enhancement for IAM regulations to support our goal for a more transparent, integer, and trustworthy assessment of IAM implementations. Thus, we add a trust score to express the quality of the auditee' IAM and the auditors' verification. We are securing this architecture with blockchain technologies to ensure integrity and availability. Sections 4.2 and 4.3 detail this enhancement of IAM audit architectures to express more trust and transparency.

Finally, we evaluate our IAM trust and reputation system in Section 5. We empirically show that the architecture can be implemented within the performance requirements of typical IAM regulations, like eIDAS, SOX, or ISO27001. Also, we analytically inspect potential attacks on trust and reputation systems and highlight their prevention or mitigation. A third evaluation examines a use case for our approach. As a wrap-up, Section 6 discusses advanced implications of this architecture, like making the trust score accessible to end users, alternatives to calculate the trust score and to execute the audit process, or its application outside of IAM.

4 IAM TRUST AND REPUTATION SYSTEM

In this Section, we design the trust and reputation system for compliance with IAM regulations. Thus, we research a classification of current IAM regulations to extract common patterns. We improve these patterns by adding computational trust for auditors and auditees, and security considerations.

[2]https://www.google.com/
[3]https://www.bing.com/

4.1 Classifying IAM Regulations

The result set of our non-exhaustive search on IAM regulations yields 52 items. We stopped querying after the realization and verification of the following phenomena. The authorities demanding regulations include (inter-)national associations and national or federal governments across the globe. The application domain of these regulations encompasses security requirements for data protection, banking, insurance, financial services, critical infrastructures, governmental or military suppliers, and general security recommendations. Overall, we classify the found IAM regulations in four distinct categories based on organizations' mandatory or voluntary participation and periodic verification processes' presence. Table 1 depicts the four resulting categories in the following paragraphs.

	(Periodic) Verification	**Without Verification**
Mandatory	Audited Compliance	Baseline Requirements
Voluntarily	Certifications	Recommendations

Table 1: Categorization of IAM Regulations.

➤ **Audited Compliance**: classifies regulations with mandatory participation and (periodic) verification processes. A regulative authority usually sets high-level security requirements, like state-of-the-art security controls or minimizing security risks. Trusted auditors verify the organizations' compliance with these requirements and, therefore, effectively interpret their details. Regulations for large organizations in the most relevant sectors opt for this type since it enforces the demanded requirements but generates expenses for external audits, limiting its feasibility for smaller organizations. Consequences for not complying with this regulation type usually include market exclusions under the regulators' authority or financial penalties. Prominent examples of this type of regulation are electronic IDentification, Authentication and trust Services (eIDAS) [13], Sarbanes-Oxley Act (SOX) [52], Health Insurance Portability and Accountability Act (HIPAA) [51], Payment Card Industry Data Security Standard (PCI-DSS) [40], North American Electric Reliability Corporation Critical Infrastructure Protection (NERC CIP) [36], or Basel III [3].

➤ **Baseline Requirements**: depict regulations that are mandatory to follow but do not impose an explicit verification process. Regulative authorities set similar security requirements for organizations like for the *Audited Compliance* type but save expenses for extensive verification processes. Due to these savings, this category allows for an application even for small organizations. Consequences for not complying with this regulation type include financial penalties or similar sanctions. Prominent examples of this type are worldwide data protection regulations on various levels[4], like the international European General Data Protection Regulation (GDPR) [12] (worldwide pioneer for many modern national data protection laws), the national Brazilian General Data Protection Law (LGPD) [10], or the federal California Consumer Privacy Act (CCPA) [8].

➤ **Certifications**: require (periodic) verification, but organizations can voluntarily obtain them. These obtained certifications provide value by setting internalized security practices on display for other

[4]A global law firm (DLA Piper) lists 163 nations with data protection laws in 2024. See world map at https://www.dlapiperdataprotection.com/index.html

parties, like end users, suppliers, or customers. The regulative authorities define the requirements for awarding their certificate, which is (periodically) verified by trusted auditors. Because of its voluntary nature, small organizations can obtain the certificate if it is relevant for them, while larger organizations might expect compliance with certain certificates from their partners. Consequences for not complying with this regulation type include not granting or terminating a certificate, which can lead to issues for business relationships if partners expect or contractually demand certain certificates. Prominent examples of this regulation type are the International Organization for Standardization (ISO) 27001 [24], System and Organization Controls 2 (SOC 2) [2], or Health Information Trust Alliance Common Security Framework (HITRUST CSF) [21].

➤ **Recommendations**: neither include a (periodic) verification process nor a mandatory participation. Regulative authorities utilize this regulation type to train or inform organizations about cyber security best practices, often in great detail. Since these informative recommendations do not need a comprehensive implementation, a recommendation effectively lowers the entry barrier for security gains for any organization. Not complying with recommendations does not impose any consequences. Prominent examples of this regulation type for IAM are the National Institute of Standards and Technology (NIST) SP 800-63 Digital Identity Guidelines [20], the UK National Cyber Security Centre (NCSC) Cloud Security Principles [35], or the Open Web Application Security Project (OWASP) Top 10 [39].

While regulations without a verification process achieve widespread inclusion or low entry barriers, large or critical organizations must prove their compliance with mandatory and periodic verification processes. Overall, this verification process includes the regulative authority (regulator) setting the cyber security requirements and audit frequency, the audited organization (auditee) obliged to comply, and the trusted auditor executing the verification. For a failed verification, the organization faces penalties, like losing the certificate or financial penalties. Especially, regulations with a (periodic) verification process, like *Audited Compliance* and *Certifications*, are relevant to this work.

4.2 Expressing Trust by Design

Utilizing the classification of IAM regulations, we synthesize four primary actors relevant for expressing trust: the regulator, the auditor, the auditee, and the public. A regulator sets the requirements that the auditee needs to follow. Auditors verify that an auditee fulfills these requirements. The public can consume the services offered by the auditee. However, the public needs to trust that organizations (auditees) will not be available if they are not secure.

We improved this synthesized design by adding a trust score so the auditee and auditor can express trust. This trust score is available to the public, while the audit details remain protected from unauthorized access. Therefore, the public can access both the trustworthiness of the auditee's implementation and the auditor's verification quality. Figure 1 depicts the essential parts of this design and its information flow. The following paragraphs detail these.

First, the regulator sets up the trust and reputation system. (1) This includes defining the requirements by appointing auditees to present an audit until a deadline. (2) The regulator also authorizes a set of trusted auditors that the auditee can hire for the audit.

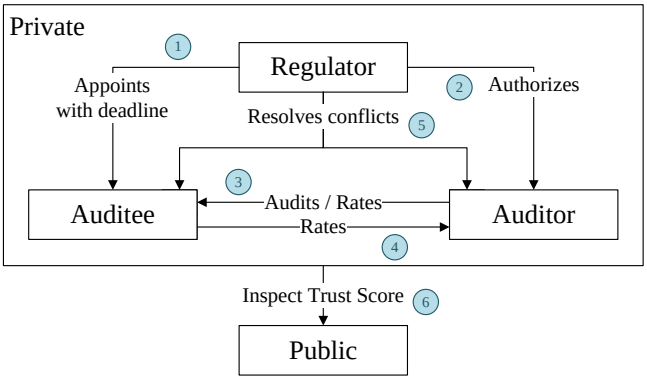

Figure 1: IAM Trust and Reputation System.

(3) The auditor then inspects the auditee. Therefore, the auditor needs access to the auditee's relevant data to base the resulting audit report on enough evidence. To this step, we add the trust score to express whether the auditee actually just complies with the bare minimum or excels. Therefore, the auditor rewards a high-performing auditee with a higher trust score, while a sufficiently performing auditee still gets a baseline score awarded. (4) Furthermore, we add a loopback rating that the auditee gives to the auditor. The loopback rating expresses the auditee's trust in the auditor regarding the proper execution of the verification process. On the one hand, this loopback trust score motivates a profound verification. On the other hand, the auditee also becomes partly responsible for the auditor's verification quality, as a highly rated but low-quality verification process deceives further (potential) auditees of the rated auditor. Additionally, this rating also sets up peer pressure among the auditees of this auditor since a failure of the verification process (c.f. Wirecard, Enron, etc.) makes past good auditor ratings from peer auditees suspicious. (5) Besides the trust rating between the auditor and auditee, the regulator resolves conflicts in exceptional cases by adjusting the ratings or demanding a second verification (dual audit). This conflict management becomes necessary upon controversies between auditors and auditees (e.g., an auditee badmouths an auditor because of a previous bad rating) or regulation violations, which should be noticed during the verification process.

Due to steps (1) to (5), the proposed trust and reputation system logs trust scores for the auditee and the auditor. So far, the system privately stores the generated trust scores and the accompanying data. This protects the sensible data of all actors from unauthorized access. (6) However, to maximize our trust score's impact, we propose making it publicly available. The public thus can access an aggregated trust score for the auditee and auditor without leaking sensible data or weaknesses.

The processed data by the trust and reputation system concludes in the data model depicted in Figure 2. In this model, the auditee and auditor inherit from the actor entity, which holds its name and an aggregated trust score for convenience. The regulator creates an instance of the audit entity for the verification process

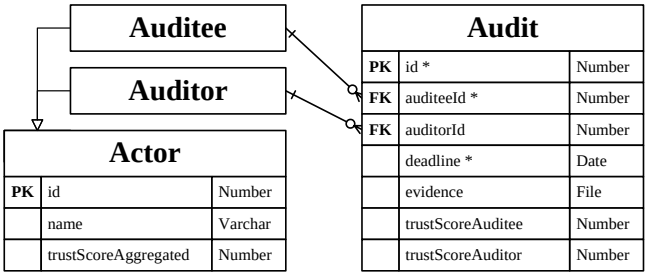

Figure 2: Simplified Data Model.

with the mandatory deadline and auditee attributes. The regulator also sets the auditor hired by the auditee. The auditee sets the trustScoreAuditor attribute, while the auditor sets the evidence and trustScoreAuditee attributes. Upon conflicts, the regulator might adjust the audit's trust scores. We expect the auditor to render the evidence to one large file containing multiple files. Additionally, a transaction history with changes and timestamps must be available for each entity, which is not explicitly depicted in Figure 2.

4.3 Securing the Trust and Reputation System

We must secure our trust and reputation system properly to ensure effective communication. Fraga et al. [15] name a few considerations for this task, including authentication, authorization, availability, and integrity. The trust and reputation system thus requires protections for unauthenticated and unauthorized access according to Section 4.2 from any user, making every user identifiable and any transaction non-reputable. The system also requires good availability to resist peak workload and ensure the necessary performance for the scope of the IAM regulations. The trust and reputation system processes also require some considerations to avoid typical attacks, for which Section 5.2 further details. We also consider the integrity of historical backup states as crucial, as auditors and auditees sometimes tend to lose essential documents [50]. While various techniques and frameworks can help us realize our trust and reputation system, an approach using blockchain technologies stands out as many properties are already included.

Wüst and Gervais [53] recommend a private-permissioned blockchain system for our trust and reputation system: we store our data model (see Figure 2) for all actors without requiring a trusted third party to host the system. All actors are identifiable and need to sign their transactions to ensure non-repudiation. Since the managed data is highly sensitive for the auditees, we omit public verifiability, leading us to a private-permissioned blockchain. Since we measure the actors' trust, we can delegate the hosting of the blockchain nodes to the most trusted actors. Therefore, the *private* parts of our proposed trust framework depicted in Figure 1 run on a private-permissioned blockchain using a smart contract for the desired logic and authorization. Only inspecting the auditee's and auditor's overall trust score is accessible to the *public*.

Blockchain technologies make past system states available, allowing for inspecting past states out of the box. Since querying the trust score for auditees and auditors is a simple read operation, the query is cachable and returns with high performance, reaching good availability. However, we recommend storing the *evidence* file in an additional off-chain storage outside of the blockchain and storing its hash value on-chain to ensure its integrity. Furthermore, this evidence file requires encryption with a state-of-the-art algorithm, like AES-256, as it contains confidential auditee data. To ensure its availability, the evidence file needs to be made accessible to the auditee, auditor, and regulator involved.

5 EVALUATION

We evaluate our trust and reputation system for IAM regulations from an empirical, analytical, and use case perspective. On the one hand, we realize and implement our design for the empirical evaluation to show that it can handle the expected workload of typical IAM regulations. Section 5.1 demonstrates that a 5-year-old commercial available notebook can run our prototype even for popular regulations. We publish our prototype as open-source on GitHub[5] so that future research can use and extend it. On the other hand, we analyze typical attacks on trust and reputation systems in Section 5.2. Finally, we present a use case to demonstrate the system's feasibility in Section 5.3.

5.1 Empirical Performance Evaluation

To demonstrate the performance necessary for our IAM trust and reputation system, we first implement a prototype, including three components: a *blockchain*, an *off-chain storage*, and a *controller*. Afterward, we evaluate the performance of our design by load tests.

➤ **A prototype for our IAM trust and reputation system.**
A *blockchain* ensures the integrity and traceability of past audit processes. For the implementation, we use the open-source blockchain framework Hyperledger Fabric[6], the best-known and most widely used private-permissioned blockchain. In contrast to other blockchains, which use pseudonymous identities, Hyperledger Fabric provides a certificate authority (CA) to ensure identifiable network participants [14]. The regulator manages this CA in our system and issues new auditors and auditees with corresponding certificates for their identity. We developed a smart contract, known as chaincode in Hyperledger Fabric, to implement the processes described in Section 4.2. Table 2 illustrates the various executable functions and the different actors that can use them. It also specifies which arguments are passed to the respective function. If a participant wishes to execute a function, they must sign the transaction for this execution in the blockchain with the private key belonging to their network identity. The chaincode then checks whether this signature matches one of the certificates issued by the CA and whether this requesting actor is authorized to use the function. For example, an auditee must not execute the "updateAuditeeScore" function and thus set its score for an audit.

As all participants in the network can view data stored in a blockchain, it is crucial to protect the actual audit data by an *off-chain storage*. This sensitive audit data needs to be encrypted and only be visible to the responsible auditor, the auditee, and the regulator. Therefore, the actors should each use their own locally operated off-chain database, which only contains the data relevant to the

[5] https://github.com/TrustInRegulations/Code
[6] https://www.hyperledger.org/

Function	Actor	Arguments
createAuditor	Regulator	name
createAuditee	Regulator	name
createAudit	Regulator	auditeeId, deadline
updateAuditor	Regulator	auditId, auditorId
updateEvidence	Regulator, Auditor	auditId, evidence
updateAuditeeScore	Regulator, Auditor	auditId, auditeeScore
updateAuditorScore	Regulator, Auditee	auditId, auditorScore
getAllAudits	Regulator, Auditor, Auditee	-
getAudit	Regulator, Auditor, Auditee	auditId
getAuditHistory	Regulator, Auditor, Auditee	auditId
getAuditorScore	Public	auditorId
getAuditeeScore	Public	auditeeId

Table 2: Chaincode Functions.

actor. We used a Couchbase[7] database as an example. However, the actors can also use their internal databases or file systems.

The *controller* is the interface for users to interact with the system. It offers a Representational State Transfer (REST) Application Programming Interface (API) and can provide a front end. Therefore, the controller receives the user input and communicates with the on-chain storage by calling the chaincode functions (Table 2) and off-chain storage. In particular, the controller encrypts and hashes the evidence file to securely store the hash value on-chain ("updateEvidence" function) and the encrypted file off-chain. Each actor should operate a local controller. This means that the actor's private key, which is used to sign transactions and prove the actor's identity, can also be stored there. Only the *Public* actor does not need to host anything. They can query the trust scores based on the controllers of the other three actors via a publicly accessible endpoint. Here, it is necessary to query different controllers and, therefore, different blockchain nodes to prevent incorrect values. We implemented the controller with Node.js[8].

➤ **Performance evaluation with load tests.**
For performance evaluation, we launched a test environment. The blockchain network is based on the open-source *test-network* configuration from Hyperledger. We used three blockchain nodes representing the three participants: regulator, auditor, and auditee. We also set up an off-chain database and a controller and deployed all components with Docker[9]. The underlying hardware is a 5-year-old commercially available notebook[10]. For measurement, we used Postman[11], which specializes in API development and testing.

For the load test, we consider a simple audit process (see Section 4.2). Thus, we executed five different chain code functions (see Table 2). The regulator creates a new audit (createAudit) and authorizes the hired auditor (updateAuditor). The auditor stores its evidence (updateEvidence) and the auditee's trust score (updateAuditeeScore), and the auditee rates the auditor back (updateAuditorScore). As a result, we simulated clients and made requests to the five corresponding REST-API endpoints of the controller, which were then forwarded to the blockchain network.

[7]https://www.couchbase.com/
[8]https://nodejs.org/en
[9]https://www.docker.com/
[10]Intel(R) Core(TM) i5-7200U 2.5 GHz with 8GB of RAM and 2 cores
[11]https://www.postman.com/

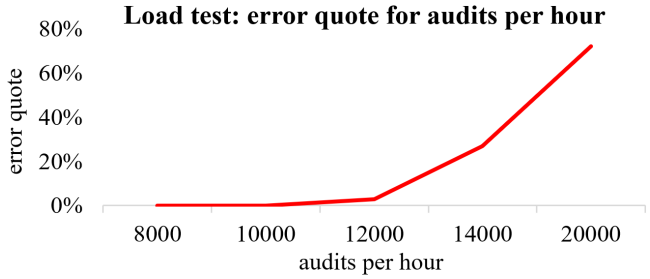

Figure 3: Error-Free Load Test for 10k Audits per Hour.

Figure 3 demonstrates the robustness of our system. It successfully processed up to 10,000 completed audits per hour without any errors. However, the error rate increased as the load increased, leading to some requests being aborted. This data showcases the system's reliability under high loads.

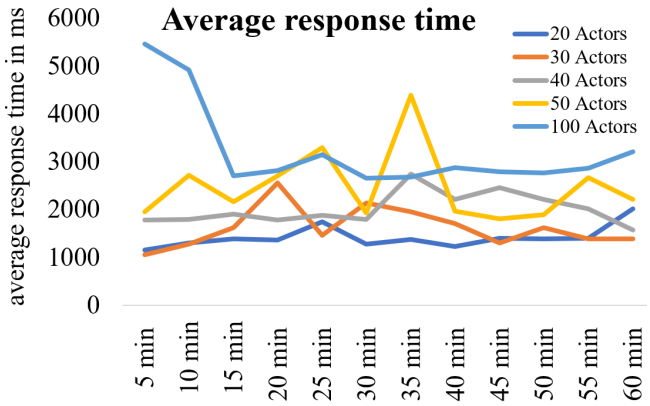

Figure 4: Average Response Times for Concurrent Actors.

Figure 4 shows the average response time for processing a single API call. For this purpose, we aggregated the average processing times of all five API endpoints. We simulated a high load by 20, 30, 40, 50, and 100 concurrent writing actors. For example, 30 actors continuously submit 30 concurrent requests over the evaluated hour, distributed across the five relevant API endpoints. The results in Figure 4 show that the response times increase with more concurrent writing actors. However, the average processing time for a single API write request is way under 10s, which is acceptable for a backend function.

$$u = \frac{f * a}{c} \tag{1}$$

Finally, to evaluate the practical applicability of our system, we related the measured performance to specific IAM regulations. Formula 1 shows the utilization percentage of the system u, which is generated by a particular IAM regulation. In the numerator, a represents the number of auditees that must be audited under this regulation, and f is the frequency of auditing. The denominator represents our trust and reputation system's total capacity c. Table 3 shows the utilization calculation for three well-known regulations. For example, SOX includes 6,000 auditees and requires an

annual audit. This results in 6,000 audits per year for SOX ($f * a$). For the total capacity of our system c, our experimental evaluation (Figure 3) showed error-free results of up to 10,000 audits per hour. With 8,766 hours per year, this results in a total capacity of $c = 87,660,000$ audits per year. Therefore, SOX utilizes our trust and reputation system by 0.0068%.

Regulation	Auditees	Frequency	Annual ($f * a$)	Utilization u
eIDAS	250	biennial	125	.0001%
SOX	6,000	annual	6,000	.0068%
ISO 27001	40,000	triennial	13,333	.0152%
Sum Σ	**46,250**	**n.a.**	**19,458**	**.0221%**

Table 3: Assumed Workload for IAM Regulations.

Also, the sum of the three well-known regulations only results in a low utilization of our system of 0.0221%. This indicates that future research might extend our IAM trust and reputation system to support multiple regulations. This evaluation, therefore, shows that our proposed trust and reputation system can process a large number of IAM regulations at the same time. This demonstrates practical applicability and availability during peak loads.

5.2 Analytical Threat Model Evaluation

During the design of the trust and reputation system, we considered threat modeling to identify potential attack vectors and vulnerabilities early. The threat analysis is based on well-known literature [15, 23, 30, 47] on attacks and defenses in online trust and reputation systems. We evaluate relevant attacks and their countermeasures in our system below.

➤ **False Information Attack**: This attack can be differentiated into *promoting* and *bad-mouthing* [23, 47]. In a promoting attack, the attacker attempts to improve a reputation in the system dishonestly. The attacker can increase its reputation (self-promoting) or another participant's. To do this, the attacker fabricates dishonest positive ratings and thus spreads false information. In bad-mouthing, the attacker creates unjustified bad feedback for another participant and seeks to damage their reputation. To prevent this attack, a trust and reputation system must ensure the authentication of the data or the sender of the rating [23]. In addition, proof of interaction should be recorded in the system before a rating is submitted [23]. This proof ensures that a rating cannot be fabricated and submitted but is linked to a previous interaction between two participants. Our reputation system authenticates the participants when they log in and stores all data generated during the audit with a non-repudiation guarantee. This audit data represents the proof of interaction. For instance, if the auditor gives the auditee an unreasonably poor rating, the auditee can check this with the help of the regulator and the recorded proof and have the rating corrected. Other auditors who have not audited this auditee cannot give a rating in the system due to a lack of interaction.

➤ **Collusion / Ballot Stuffing Attack**: In a collusion attack, various participants collaborate to conduct malicious actions in the system [30, 47]. The ballot stuffing attack is particularly relevant, where some participants inflate the system with good ratings while actually showing bad behavior [30]. This means that an auditee may not pass its audit, but the auditor and auditee still give each other a good rating. Several opinions from different participants can be requested for a new rating [30] to counteract this. Joint and dual audits, in which several auditors are responsible for auditing an auditee, are suitable for this purpose (see Section 6). For example, when suspicion about Wirecard and EY rose, KPMG was tasked with an ad-hoc dual audit, revealing the contradictions [50]. In addition, the regulator can rotate the assigned auditors every few years. Due to the integrity and non-repudiation of the stored audit data ensured by the blockchain, the new auditors would notice any past collusion attacks. They can understand the entire audit history and access the off-chain storage. The regulator can initiate this.

➤ **Sybil Attack**: For this attack, an attacker creates various, usually pseudonymous, identities in the system, which are used for malicious activities [23, 30]. This allows an attacker to submit various false information, which can be used for attacks such as promotion or bad-mouthing, as described above. However, our trust and reputation system is based on a private-permissioned blockchain in which the system identities are linked to real-world identities. In addition, only the regulator issues an identity for a new participant.

➤ **Re-Entry Attack**: A participant with a poor reputation can abandon the associated identity and re-enter the system with a new unrated identity [15, 23]. If successful, the two identities cannot be linked to each other. For example, an auditee that has received bad ratings from auditors could create a new user account. However, the attack is not possible because the regulator assigns the identities in the system. This attack is related to the sybil attack. In a re-entry attack, only one active identity is used at a time, while the attacker uses several identities simultaneously in the Sybil attack [30].

➤ **Traitor Attack**: In this attack, the attacker first builds up a good reputation over a long period through good behavior. The attacker can then misbehave for a short time. The long history of good ratings will outweigh the recent actions so that the reputation value remains almost unchanged [23]. For example, an auditor has built up a good reputation with many smaller and less complex audits. In an expensive audit, the auditor saves resources and receives a lower rating from the auditee. To counter this attack, it is helpful to reduce the weighting of past assessments when calculating the overall reputation score (see Section 6) [30, 47]. Furthermore, current actions are weighted more heavily in the calculation. This makes bad behavior noticeable in the system more quickly.

➤ **Data Manipulation Attack**: Individual trusted actors who are responsible for the calculation and distribution of reputation values pose a risk to the system [23]. If they decide to act maliciously or are compromised by an attack, this can cause significant damage to participants unnoticed. These actors can manipulate the submitted reputation data and the value aggregated from it, which is then disseminated in the system [15, 30]. Since the blockchain decentralizes all references to audit and reputation data while storing them tamper-proof, this attack is practically impossible in our system.

➤ **Denial of Service Attack**: In this attack, the attacker puts the system out of operation by overloading it [23, 47]. Central systems without redundancy are particularly vulnerable to this [23]. This means that trust score submissions and reputation queries are not available. As in our approach, distributed storage and reputation calculation can resist this attack.

5.3 Use Case Evaluation

We present a realistic use case with an exemplary auditee (Bank B) to demonstrate the feasibility of our proposed system.

Bank *B* operating in the US financial market is subject to SOX regulation and hires Auditor *A1* for the audit. According to SOX SEC.302 (a)(4)(A), B is *"responsible for establishing and maintaining internal controls"* [52]. During the audit, A1 found that B had documentation issues. SEC.302 (a)(5)(A) requires documentation and audit for *"all significant deficiencies in the design or operation of internal controls"* [52]. According to A1, this disclosure by B was only partially sufficient. On a trust score scale of 1 - 10, A1 gives B a score of 5. Based on the completed audit report, Bank B partially recognizes the documentation issue but assumes that A1 is insufficiently familiar with B's internal controls. The transparent history of other auditees' trust scores for A1 shows a declining rating. B thus requests a second opinion from the SOX regulator, the U.S. Securities and Exchange Commission (SEC). The SEC investigates the concerns expressed and allows a second Auditor *A2* to conduct the audit. After its assessment, A2 also recognized weaknesses in B's internal controls, but A2 awarded a better trust score of 7, considering A1's assessment to be too low. The arithmetic mean of the two trust score values, 5 and 7, is 6 for this auditing cycle. In the following year, B now engages A2 directly. B has remedied the existing weaknesses in internal controls and receives the best rating from A2 due to full compliance with a trust score of 10. In the third year, bank B only partially fulfills a specific paragraph in the SOX regulation. SEC.302 (a)(4)(C) requires an effectiveness evaluation of *"internal controls as of a date within 90 days prior to the report"* [52]. However, some reviews to evaluate the effectiveness were older than 90 days. As a result, A2 reduced its trust score evaluation by one point compared to the previous year and awarded a trust score of 9. After the third year, the arithmetic mean results in a total trust score of $(6 + 10 + 9)/3 = 8.33$ for B.

Overall, the public can see that B has improved its trust score over the years, but the system also shows recent deterioration with a slight downward trend. This can motivate B to regain the best rating in the following year. Past assessments on A1, including that of B, have identified it as a relatively worse auditor. Other auditees subject to the SOX regulation can thus use our proposed system to assess the trustworthiness of auditors before hiring one. For auditor A1, these assessments can motivate to audit better and more thoroughly to continue receiving audit mandates.

6 DISCUSSION

➤ **Trust Score Calculation**: Our trust and reputation system, which calculates an arithmetic mean for the trust scores of auditees and auditors, can benefit from more advanced approaches that consider further crucial aspects.

One such crucial dimension is timeliness, where the most recent audit rating could serve as the most timely trust score, albeit at the expense of disregarding past trust scores, making it vulnerable to outliers. However, it's important to note that past trust scores provide a comprehensive view of the auditees' and auditors' trustworthiness. Therefore, considering past trust scores not only adds a further layer of meaning but also ensures the system's thoroughness and reliability. While recent ratings express current trust levels, also considering previous trust scores depicts a more holistic picture. For example, by including only the three most recent trust scores, we can express a current trust assessment without sticking to an outdated past or being endangered by a recent outlier trust rating. For these recent few trust scores, we can again apply a central tendency (like the arithmetic mean) or degrade them over time. By degrading the trust scores over time, we can stress the most recent ones the most while fading out older ones. For example, we can calculate the current trust score by weighting the most recent rating with 50%, the second most recent one with 30%, and the third most recent one with 20%. Another exciting display of past trust scores is their development. Considering its development, we can express whether an auditee or auditor has improved or worsened since the last audit(s). We think these depictions of trends are helpful since an end user might consider an outdated perception of trust and thus can comprehend recent improvement or deterioration.

Besides the timeliness dimension of the trust score, another dimension for an advanced trust score calculation is the relationship between auditees and auditors. The primary idea for considering their relationship is that the same entities that rate a trust score also receive one. For example, if an auditor rates higher trust scores while it receives lower ones from its auditees, this auditor's trust scores have dubious validity. To resolve this situation, the regulator can intervene and decrease the weighting of these audits, decrease higher (or increase lower) trust scores from this auditor, regress this auditor's ratings towards a general central tendency, request another audit (dual audit), or remove this auditor from the set of trusted auditors. One advantageous perspective on the actors' relationships within the reputation system is benchmarking similar actors. By grouping actors based on their services or sectors, we can effectively express their trust in relation to their peers. This approach creates a sense of peer pressure for these actors and presents more trusted alternatives for the end-users.

While this paragraph outlines a few approaches for advanced utilization of the trust score within the reputation system alone, it's important to note that external data can also play a significant role. This external data, including recent events for auditees and auditors, introduces additional dependencies for the reputation system. Overall, the trust score calculation can express further details and focus on the trust and reputation system.

➤ **Joint and Dual Audits**: While one audit usually includes one auditor and one auditee, joint or dual audits can improve the quality of audits by adding more auditors (and audits). Thus, these quality improvements also affect the trust score. Our trust and reputation system can model both by adding relationships with a Many-To-Many cardinality between two audit entities. On the one hand, for a *Joint Audit*, two auditors create a single audit report together. For a joint audit, we suggest creating an audit for each auditor and connecting them via a relationship between the joint audits. The hash value for the evidence file will be the same as the auditors base their rating on the same audit report and evidence. However, by representing the joint audit as two connected audits, each auditor rates the trust score for the auditee independently. Vice versa, the auditee rates the auditors independently. This approach allows for a finer expression of trust for each actor. On the other hand, for a *Dual Audit*, two auditors create two audit reports independently. Similar to a joint audit, a relationship between each conducted audit

indicates their duality. Also, we can express the trust scores for each audit independently. Unlike the joint audit, the evidence file and, thus, its hash value will differ since the auditors will create different audit reports concurrently. A regulator might also set up a dual audit if a second opinion seems necessary. Nevertheless, trust score calculations must adapt when applying these extensions regardless of a joint or dual audit because the increased number of audits ballot-stuff previous ratings. For example, an auditee might request more audits to hide a poor audit rating from the first auditor. Combining the trust ratings with a central tendency (like the arithmetic mean) reduces the audit number for an even comparison. Overall, applying additional audits strengthens the quality of the trust score, similar to the four-eye principle. Our proposed reputation system for expressing trust in IAM regulations can model these by adding a relationship between the audits.

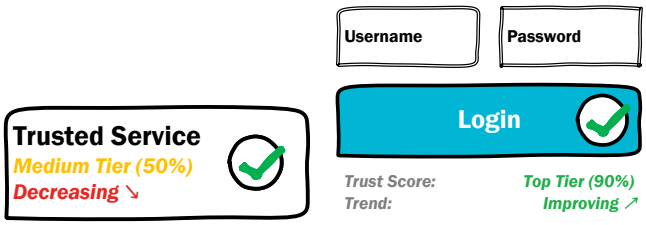

(a) Badge (b) Login

Figure 5: Mock-ups for an Exemplary Trust Score Display.

➤ **Reputation Badge**: While rating and calculating the trust score is important, communicating the trust score to "the public" end-user is one of the most crucial parts of our trust and reputation system. As described earlier, we can express the trust score as a central tendency (like the arithmetic mean), inspect its trend (like an increase or decrease), or show it relative to its peers. Therefore, Figure 5 displays two exemplary mock-ups for presenting the trust score on crucial steps for user experience: as a badge for the website of an auditee (see Figure 5a) and an extension for a login form of the auditee's service (see Figure 5b). The percentages display the trust score itself, the trend shows whether the last audits improved or worsened the trust score, and the tiers depict the relation of the auditee's trust score with its peers. The badge in Figure 5a shows an auditee who slacked off but is still recognized as a trusted service. Figure 5b displays a service login that earned high trust. In both exemplary mock-ups, a public end-user can easily comprehend the trustworthiness of the auditee's services.

➤ **Not-Limited to IAM or a Single Regulation**: We considered regulations (in-)directly demanding IAM only. However, the need for internal controls or proper controls for authentication and authorization is very common, as seen by the broad application domain of the found IAM regulations. While our trust and reputation system bases on IAM regulations, the classification of regulations by participation & verification and the resulting reputation system should apply to any regulation. Therefore, expanding our reputation system to include any regulation that includes a (periodic) verification process might be possible. Furthermore, as seen in the evaluation Section 5, future work might extend our reputation system to cover several regulations with several regulators. This high coverage of regulations also improves the expressiveness of trust and streamlines the trust score for public end-users.

7 CONCLUSION

Regulations are a common theme and motivator for IAM [28], while their fulfillment is handled rather silently for the public end-user. Therefore, we designed a trust and reputation based on IAM regulations. This includes a classification of current and prominent IAM regulations, designing the trust and reputation system, and securing it with private-permissioned blockchain technologies to ensure privacy for auditees, tamper-proof storage of the trust score, and substantial availability due to decentralization. We have shown sufficient performance for several regulations, protections for typical attacks on trust and reputation systems, and feasibility for a use case. Furthermore, we discuss the calculation of the trust score, extensions for the verification process, communication of the trust score for maximizing its impact for the public end-user, and its application potential outside of IAM. For future work, we propose to extend the trust and reputation system to support several regulations simultaneously or to study the usability of the trust score for an effective presentation to the public end-users.

ACKNOWLEDGMENTS

The German Federal Ministry of Education and Research supported the research leading to these results as part of the DEFENSIVE project (https://defensive.it-sicherheitscluster.de/).

REFERENCES

[1] Muhammad Aaqib, Aftab Ali, Liming Chen, and Omar Nibouche. 2023. IoT trust and reputation: a survey and taxonomy. *Journal of Cloud Computing* 12, 1 (22 03 2023), 42. https://doi.org/10.1186/s13677-023-00416-8

[2] American Institute of Certified Public Accountants. 2022. SOC 2® Reporting on an Examination of Controls at a Service Organization Relevant to Security, Availability, Processing Integrity, Confidentiality, or Privacy. https://www.aicpa-cima.com/cpe-learning/publication/soc-2-reporting-on-an-examination-of-controls-at-a-service-organization-relevant-to-security-availability-processing-integrity-confidentiality-or-privacy

[3] Basel Committee on Banking Supervision. 2011. Basel III: A global regulatory framework for more resilient banks and banking systems.

[4] 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

[5] Rafael Belchior, André Vasconcelos, Sérgio Guerreiro, and Miguel Correia. 2021. A Survey on Blockchain Interoperability: Past, Present, and Future Trends. *ACM Comput. Surv.* 54, 8, Article 168 (oct 2021), 41 pages. https://doi.org/10.1145/3471140

[6] Emanuele Bellini, Youssef Iraqi, and Ernesto Damiani. 2020. Blockchain-Based Distributed Trust and Reputation Management Systems: A Survey. *IEEE Access* 8 (2020), 21127–21151. https://doi.org/10.1109/ACCESS.2020.2969820

[7] Diego De Siqueira Braga, Marco Niemann, Bernd Hellingrath, and Fernando Buarque De Lima Neto. 2018. Survey on Computational Trust and Reputation Models. *ACM Comput. Surv.* 51, 5, Article 101 (11 2018), 40 pages. https://doi.org/10.1145/3236008

[8] California State Legislature. 2018. California Consumer Privacy Act. https://oag.ca.gov/privacy/ccpa

[9] Yuan Cao and Lin Yang. 2010. A survey of Identity Management technology. In *2010 IEEE International Conference on Information Theory and Information Security*. IEEE, Beijing, China, 287–293. https://doi.org/10.1109/ICITIS.2010.5689468

[10] Câmara dos Deputados. 2018. Lei Geral de Proteção de Dados Pessoais. https://www.planalto.gov.br/ccivil_03/_Ato2015-2018/2018/Lei/L13709.htm

[11] Nabil El Ioini and Claus Pahl. 2018. A Review of Distributed Ledger Technologies. In *On the Move to Meaningful Internet Systems. OTM 2018 Conferences*, Hervé Panetto, Christophe Debruyne, Henderik A. Proper, Claudio Agostino Ardagna, Dumitru Roman, and Robert Meersman (Eds.). Springer International Publishing, Cham, 277–288.

[12] European Commission. 2016. General Data Protection Regulation.

[13] European Parliament, Council of the European Union. 2014. Regulation (EU) No 910/2014 of the European Parliament and of the Council of 23 July 2014 on electronic identification and trust services for electronic transactions in the internal market and repealing Directive 1999/93/EC.

[14] Hyperledger Foundation. 2024. Hyperledger Fabric Documentation. https://hyperledger-fabric.readthedocs.io/en/release-2.5/index.html

[15] D. Fraga, Z. Bankovic, and J.M. Moya. 2012. A Taxonomy of Trust and Reputation System Attacks. In *2012 IEEE 11th International Conference on Trust, Security and Privacy in Computing and Communications*. IEEE, Liverpool, UK, 41–50. https://doi.org/10.1109/TrustCom.2012.58

[16] Ludwig Fuchs and Gunther Pernul. 2007. 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)*. IEEE, Vienna, Austria, 374–384. https://doi.org/10.1109/ARES.2007.145

[17] Ludwig Fuchs, Günther Pernul, and Ravi Sandhu. 2011. Roles in information security – A survey and classification of the research area. *Computers & Security* 30, 8 (2011), 748–769. https://doi.org/10.1016/j.cose.2011.08.002

[18] Thomas Geras and Thomas Schreck. 2023. Sharing Communities: The Good, the Bad, and the Ugly. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security* (Copenhagen, Denmark) *(CCS '23)*. Association for Computing Machinery, New York, NY, USA, 2755–2769. https://doi.org/10.1145/3576915.3623144

[19] Jones Granatyr, Vanderson Botelho, Otto Robert Lessing, Edson Emílio Scalabrin, Jean-Paul Barthès, and Fabrício Enembreck. 2015. Trust and Reputation Models for Multiagent Systems. *ACM Comput. Surv.* 48, 2, Article 27 (10 2015), 42 pages. https://doi.org/10.1145/2816826

[20] 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

[21] Health Information Trust Alliance. 2023. Health Information Trust Alliance Common Security Framework. https://hitrustalliance.net/hitrust-framework

[22] Paul M. Healy and Krishna G. Palepu. 2003. The Fall of Enron. *Journal of Economic Perspectives* 17, 2 (06 2003), 3–26. https://doi.org/10.1257/089533003765888403

[23] Kevin J. Hoffman, David Zage, and Cristina Nita-Rotaru. 2009. A survey of attack and defense techniques for reputation systems. *ACM Comput. Surv.* 42, 1 (2009), 1:1–1:31. https://doi.org/10.1145/1592451.1592452

[24] ISO 27001. 2013. *Information technology — Security techniques — Information security management systems — Requirements.* Standard. International Organization for Standardization.

[25] Arthur S. Jacobs, Roman Beltiukov, Walter Willinger, Ronaldo A. Ferreira, Arpit Gupta, and Lisandro Z. Granville. 2022. AI/ML for Network Security: The Emperor has no Clothes. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security* (Los Angeles, CA, USA) *(CCS '22)*. Association for Computing Machinery, New York, NY, USA, 1537–1551. https://doi.org/10.1145/3548606.3560609

[26] Audun Jøsang, Roslan Ismail, and Colin Boyd. 2007. A survey of trust and reputation systems for online service provision. *Decision Support Systems* 43, 2 (2007), 618–644. https://doi.org/10.1016/j.dss.2005.05.019 Emerging Issues in Collaborative Commerce.

[27] Michael Kapoor. 2023. Deloitte Canada Fined $1.1 Million for Backdating Audit Papers. https://news.bloombergtax.com/financial-accounting/deloitte-canada-fined-1-1-million-for-backdating-audit-papers

[28] Sascha Kern, Thomas Baumer, Ludwig Fuchs, and Günther Pernul. 2023. Maintain High-Quality Access Control Policies: An Academic and Practice-Driven Approach. In *Data and Applications Security and Privacy XXXVII*, Vijayalakshmi Atluri and Anna Lisa Ferrara (Eds.). Springer Nature Switzerland, Cham, 223–242. https://doi.org/10.1007/978-3-031-37586-6_14

[29] Sascha Kern, Thomas Baumer, Sebastian Groll, Ludwig Fuchs, and Günther Pernul. 2022. Optimization of Access Control Policies. *Journal of Information Security and Applications* 70 (2022), 103301. https://doi.org/10.1016/j.jisa.2022.103301

[30] Eleni Koutrouli and Aphrodite Tsalgatidou. 2012. Taxonomy of attacks and defense mechanisms in P2P reputation systems - Lessons for reputation system designers. *Comput. Sci. Rev.* 6, 2-3 (2012), 47–70. https://doi.org/10.1016/J.COSREV.2012.01.002

[31] Yong Kuang, Hongyun Xu, Rui Jiang, and Zhikang Liu. 2022. GTMS: A Gated Linear Unit Based Trust Management System for Internet of Vehicles Using Blockchain Technology. In *2022 IEEE International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom)*. IEEE, Wuhan, China, 28–35. https://doi.org/10.1109/TrustCom56396.2022.00015

[32] Chi-Wei Lien and Sudip Vhaduri. 2023. Challenges and Opportunities of Biometric User Authentication in the Age of IoT: A Survey. *ACM Comput. Surv.* 56, 1, Article 14 (08 2023), 37 pages. https://doi.org/10.1145/3603705

[33] Mark Maurer. 2024. KPMG Fined Record $25 Million in Exam-Cheating Scandal. https://www.wsj.com/articles/kpmg-fined-25-million-over-alleged-netherlands-exam-cheating-a4dcba2a

[34] Barsha Mitra, Shamik Sural, Jaideep Vaidya, and Vijayalakshmi Atluri. 2016. A Survey of Role Mining. *ACM Comput. Surv.* 48, 4, Article 50 (02 2016), 37 pages. https://doi.org/10.1145/2871148

[35] National Cyber Security Centre. 2018. Cloud security guidance. https://www.ncsc.gov.uk/collection/cloud

[36] North American Electric Reliability Corporation. 2024. Critical Infrastructure Protection. https://www.nerc.com/pa/Stand/Pages/ReliabilityStandards.aspx Series of standards.

[37] Martin Nuss, Alexander Puchta, and Michael Kunz. 2018. Towards Blockchain-Based Identity and Access Management for Internet of Things in Enterprises. In *Trust, Privacy and Security in Digital Business*, Steven Furnell, Haralambos Mouratidis, and Günther Pernul (Eds.). Springer International Publishing, Cham, 167–181. https://doi.org/10.1007/978-3-319-98385-1_12

[38] Aleksandr Ometov, Sergey Bezzateev, Niko Mäkitalo, Sergey Andreev, Tommi Mikkonen, and Yevgeni Koucheryavy. 2018. Multi-Factor Authentication: A Survey. *Cryptography* 2, 1 (2018), 1–31. https://doi.org/10.3390/cryptography2010001

[39] OWASP Top 10 team. 2021. OWASP Top10. https://owasp.org/Top10/

[40] PCI Security Standards Council. 2022. PCI DSS: v4.0. https://docs-prv.pcisecuritystandards.org/PCI%20DSS/Standard/PCI-DSS-v4_0.pdf

[41] J Ratnatunga. 2023. PwC Tax Scandal's Aftermath: It's Time to Seriously Regulate the Big 4. https://cmaaustralia.edu.au/ontarget/pwc-tax-scandals-aftermath-its-time-to-seriously-regulate-the-big-4/

[42] Daniël Reijsbergen, Aung Maw, Zheng Yang, Tien Tuan Anh Dinh, and Jianying Zhou. 2023. TAP: Transparent and Privacy-Preserving Data Services. In *32nd USENIX Security Symposium (USENIX Security 23)*. USENIX Association, Anaheim, CA, 6489–6506. https://www.usenix.org/conference/usenixsecurity23/presentation/reijsbergen

[43] Ravi S. Sandhu. 1998. 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 *Advances in Computers*, Marvin V. Zelkowitz (Ed.). Advances in Computers, Vol. 46. Elsevier, online, 237–286. https://doi.org/10.1016/S0065-2458(08)60206-5

[44] Sarwar Sayeed, Nikolaos Pitropakis, William J. Buchanan, Evangelos Markakis, Dimitra Papatsaroucha, and Ilias Politis. 2023. TRUSTEE: Towards the creation of secure, trustworthy and privacy-preserving framework. In *Proceedings of the 18th International Conference on Availability, Reliability and Security* (Benevento, Italy) *(ARES '23)*. Association for Computing Machinery, New York, NY, USA, Article 145, 10 pages. https://doi.org/10.1145/3600160.3604997

[45] Peiyao Sheng, Xuechao Wang, Sreeram Kannan, Kartik Nayak, and Pramod Viswanath. 2023. TrustBoost: Boosting Trust among Interoperable Blockchains. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security* (Copenhagen, Denmark) *(CCS '23)*. Association for Computing Machinery, New York, NY, USA, 1571–1584. https://doi.org/10.1145/3576915.3623080

[46] Stavros Simou, Aikaterini-Georgia Mavroeidi, and Christos Kalloniatis. 2024. Review on Privacy and Trust Methodologies in Cloud Computing. In *Computer Security. ESORICS 2023 International Workshops*, Sokratis Katsikas, Frédéric Cuppens, Nora Cuppens-Boulahia, Costas Lambrinoudakis, Joaquin Garcia-Alfaro, Guillermo Navarro-Arribas, Pantaleone Nespoli, Christos Kalloniatis, John Mylopoulos, Annie Antón, and Stefanos Gritzalis (Eds.). Springer Nature Switzerland, Cham, 494–505.

[47] Yan Sun and Yuhong Liu. 2012. Security of Online Reputation Systems: The evolution of attacks and defenses. *IEEE Signal Process. Mag.* 29, 2 (2012), 87–97. https://doi.org/10.1109/MSP.2011.942344

[48] Ali Sunyaev. 2020. Distributed Ledger Technology. In *Internet Computing*, Ali Sunyaev (Ed.). Vol. 2. Springer International Publishing, Cham, 265–299. https://doi.org/10.1007/978-3-030-34957-8_9

[49] S. Tan, W. Chen, R. Deng, and R. Popa. 2023. MPCAuth: Multi-factor Authentication for Distributed-trust Systems. In *2023 IEEE Symposium on Security and Privacy (SP)*. IEEE Computer Society, Los Alamitos, CA, USA, 829–847. https://doi.org/10.1109/SP46215.2023.10179481

[50] Fabian Maximilian Johannes Teichmann, Sonia Ruxandra Boticiu, and Bruno S. Sergi. 2023. Wirecard scandal. A commentary on the biggest accounting fraud in Germany's post-war history. *Journal of Financial Crime* ahead-of-print, ahead-of-print (01 01 2023), 1–8. https://doi.org/10.1108/JFC-12-2022-0301

[51] United States Congress. 1996. Health Insurance Portability and Accountability Act of 1996.

[52] United States Congress. 2002. Sarbanes-Oxley Act of 2002. Corporate responsibility.

[53] Karl Wüst and Arthur Gervais. 2018. Do you Need a Blockchain?. In *2018 Crypto Valley Conference on Blockchain Technology (CVCBT)*. IEEE, Zug, Switzerland, 45–54. https://doi.org/10.1109/CVCBT.2018.00011

[54] Han Yu, Zhiqi Shen, Chunyan Miao, Cyril Leung, and Dusit Niyato. 2010. A Survey of Trust and Reputation Management Systems in Wireless Communications. *Proc. IEEE* 98, 10 (10 2010), 1755–1772. https://doi.org/10.1109/JPROC.2010.2059690

Received 10 May 2024; accepted 02 June 2024