

Original Article

Automated Detection of Excessive Privileges in Enterprise Cloud Environments

Uday Allala

Lead SRE Engineer, USA.

Received: 05-07-2020

Revised: 05-28-2020

Accepted: 06-05-2020

Published: 06-09-2020

ABSTRACT

As enterprise cloud services become ubiquitous, the management of identities and access has become more complex, and it is crucial to identify excessive privileges as a security concern. The proliferation of enterprise cloud services has made identity and access management more complex, and it has become important to detect excessive privileges as a security concern. Users, service accounts, and applications often end up with more permissions than needed, which can lead to unnecessary attack surface and risk of unauthorized access, privilege escalation, and misuse of sensitive cloud resources. Traditional access control and periodic auditing methods are primarily rule-based and manual, and fail to effectively discover unused, unnecessary and high-risk privileges in dynamic cloud environments continuously. To overcome these restrictions, this paper presents an automated detection framework of excessive privileges that considers identity attributes, roles assigned, resources sensitivity, and access behavior observed to detect potentially unnecessary privileges. The methodology proposed integrates the analysis of permissions, assessment of privileges with risks, behavioral features with automated classification to continuously analyze the appropriateness of access that has been granted. Experimental evaluation of the proposed framework is done with representative enterprise cloud identity and access scenarios, and then compared with traditional privilege analysis methods. The accuracy, precision, recall, F1-score, false-positive rate, detection time, and computational overhead are used to measure performance. The evaluation shows the ability of the proposed approach to detect over-privileged and unused permissions while achieving good detection rates in different identity and permission settings. The framework also introduces a risk-based prioritization for identified privileges so that security administrators can prioritize remediation activities for high-risk access relationships. The primary idea of the work is an automated and scalable solution to detect excessive privileges continuously that enhances the least-privilege enforcement, and enhances the security posture of enterprise cloud environments by supporting cloud identity governance.

KEYWORDS

Excessive Privileges, Privilege Detection, Cloud Security, Identity and Access Management, Least Privilege, Access Control, Zero Trust, Machine Learning, Cloud-Native Security, Identity Governance.

1. INTRODUCTION

1.1. Background and Motivation

Cloud computing has revolutionized enterprise IT by allowing applications, services, data, and computing resources to be provisioned dynamically within distributed cloud computing environments. [1] This has also made identity and access management complex (IAM), as enterprises have to handle numerous users, service accounts, apps, roles and permissions. In such instances, the concept of "least privilege" is key to guaranteeing that each identity only has access to what it needs to accomplish the duties it was granted to provide. But as organizations, their resources and permission access needs evolve over time, users and applications can end up with permissions they no longer need. Given this, how enterprise cloud infrastructures can be better protected by automated analysis of identity permissions and access behavior is becoming ever more relevant, with a focus on excess privileges, unnecessary attack surfaces, and the overall security posture of enterprise cloud infrastructures.

1.2. Problem of Excessive Privileges

Excessive privileges are when users, applications, or accounts have permissions that are not needed for their normal operations. Inappropriate assignments, accumulation of privileges, unused permissions, and changes in job function and resource needs are all examples of potential causes of such permissions. In a large cloud deployment, these conditions are not easy to find manually, because permissions are spread out throughout the various services, and permissions often evolve over time. Allowing overly privileged accounts can introduce a greater risk of compromised accounts and can provide unauthorized access to valuable resources. Current methods, which focus mainly on static policies, pre-defined rules, and periodic access reviews, may not reflect the changing patterns of usage. Enterprises thus need a scalable and automated way to continuously monitor the permissions assigned to users and access patterns to accurately identify and prioritize excessive permissions within the cloud.

2. BACKGROUND AND RELATED WORK

Table 1: Comparison of Existing Privilege Detection Approaches

Approach	Privilege Analysis	Usage Behavior	Context Awareness	ML Support	Continuous Detection	Main Limitation
Static Rule-Based Detection	High	Low	Low	No	Limited	Depends on predefined rules
Role-Based Analysis	High	Low	Medium	No	Limited	May not capture actual usage
Unused Permission Analysis	Medium	High	Low	No	Moderate	May flag legitimately rare access
Anomaly-Based	Medium	High	Medium	Yes	High	Sensitive to behavioral

Detection						variations
Traditional ML Classification	Medium	High	Medium	Yes	High	Requires suitable training data
Policy-Based Privilege Analysis	High	Low	Medium	No	Moderate	Limited adaptability to dynamic environments
Proposed Framework	High	High	High	Yes	High	Requires cloud IAM and access-event data

2.1. Privilege Creep and Excessive Permissions

Privilege creep is the gradual accumulation of access rights that users, applications and service accounts obtain over time that are excessive for their job functions. In enterprise cloud deployments privilege creep can happen as a result of role changes, temporary need, permissions being inherited, organizational changes, and the constant introduction of new cloud services. [2] Identities may have access to sensitive resources that rarely or never get used, as permissions build up, which violates the principle of least privilege and expands the possible attack surface. Dynamic cloud infrastructures, which can include access policies defined across multiple services, roles, resources and accounts, are a particular challenge when it comes to identifying excessive permissions. As such, regular and thorough review of assigned permissions in combination with real-world usage is necessary to identify legitimate permissions, or permissions that are not necessary or excessive, and to enable ongoing privilege governance.

2.2. Existing Automated Privilege Detection Techniques

Current automated privilege detection methods typically focus on access policies, role mappings, permission usage, identity behaviour and resource sensitivity to identify excess permissions and risky access rights. While rule-based techniques can be used to detect known violations like unused permissions, excessive role assignments, and access to sensitive resources, behavioral approaches and machine learning can be used to detect deviations from normal access, and to estimate privilege-related risks. In many cases, though, the existing techniques rely on static policies, hard-coded thresholds, limited context information, or even periodic analysis, which can mitigate their ability to be effective in ever-changing enterprise cloud applications. Also, methods that only consider mandatory permissions could be ineffective at differentiating between the unnecessary and the necessary permissions that are needed for operational workflows. The restrictions point to the necessity of an automated framework for detecting excessive privileges, one that is context-aware and scalable, and that incorporates a mix of permission configurations, actual usage behavior, identity characteristics, and risk assessment.

3. PROPOSED AUTOMATED PRIVILEGE DETECTION FRAMEWORK



Fig 1 - Overall Architecture of the Proposed Automated Privilege Detection Framework

Table 2: Components of Proposed Automated Privilege Detection Framework

Framework Component	Primary Input	Main Processing	Output
Identity and Permission Collection	Users, applications, service accounts	Collect IAM and permission information	Identity-permission records
Role and Identity Profiling	Identity and role attributes	Construct identity and role profiles	User/role profiles
Permission Usage Analysis	Access-event logs	Analyze frequency and recency of permission usage	Usage characteristics
Contextual Privilege Analysis	Roles, resources, usage behavior	Evaluate privileges within operational context	Contextual privilege indicators
Excessive Privilege Identification	Assigned and observed permissions	Detect unnecessary or excessive access	Candidate excessive privileges
Risk Scoring	Privilege and behavioral features	Calculate normalized privilege risk	Risk score
ML-Based Classification	Constructed feature vectors	Predict privilege status/risk	Privilege classification

3.1. System Overview

The proposed APF (Automated Privilege Detection Framework) aims to automate the identification of excessive, unnecessary and potentially high-risk privileges that may be held by users, applications and service accounts in enterprise cloud environments on an ongoing basis. [3] The framework incorporates cloud identity details, permissions and role configuration, resource sensitivity, and the observed access activity to determine if the permissions granted match the normal requirements for operations. The overall process involves identity and permission collection, user and role profiling, excessive privileges identification, risk scoring and automated classification of

permissions. The framework will overcome the shortcomings of periodic and rule-based privilege review and offer a scalable way to continuously enforce least privilege.

3.2. System Architecture

The structure of the proposed framework is made up of various interconnected modules that are responsible for gathering, processing, analyzing, and reacting to privilege information from the Cloud. [4] The data collection layer collects the attributes of identities, role assignments, permissions, resource information and access-event records from a cloud environment. The processed layer normalizes the information collected and builds user, role, permission and resource profiles. The detection layer looks at permission usage and determines if actual usage is deviating from permission usage, and the risk-analysis layer looks at the sensitivity of the permission, its frequency of being used, the criticality of the resource and if there are any behavioural anomalies to assign a risk score. Last but not least, the classification and remediation layers prioritize detected excessive privileges, and facilitate appropriate access reduction or review actions.

3.3. Permission Usage Analysis

Permission usage analysis is used to determine if there is a correlation between privileges that an identity is given and privileges that are used during typical cloud operations. The framework reviews access-event data for each user, application or service account and identifies the most common, least common, oftentimes unused and potentially unnecessary access permissions. Analysis takes account of how often users use them, how recent their accesses have been, how sensitive the resources are to access, the type of operation and the correlation between permissions and assigned responsibilities. Those permissions that are still in effect but don't show signs of legitimate usage are flagged for further review, instead of being outright revoked. This usage-based analysis is able to differentiate between potentially excessive privileges and legitimate permissions for occasional or critical operational activities.

3.4. Excessive Privilege Identification

Excessive privilege identification checks if the privilege assigned is more than what an identity requires within its role, actions, and context of operation. [5] The framework assesses usage patterns, such as unused access rights, unnecessary administrative capabilities, scope of resources, access to sensitive resources, and differences in role responsibilities and assigned rights. More than one contextual characteristic is used to analyze each permission, not just prespecified access rules. Permissions that meet the excessive-privilege criteria of the framework are marked for risk assessment and classification; permissions that do not meet the criteria are kept. Such a strategy allows for the ongoing discovery of potential redundant access rights as cloud identities, roles, resources and operational requirements change over time.

3.5. Privilege Risk Scoring

The privilege risk-scoring mechanism provides a numeric risk score for each and every potential privilege that is overused, using multiple security and usage properties. [6] Permission

sensitivity, Resource Criticality, Privilege Scope, Access Frequency, Access recency, Identity Type, Behavioral deviation are used as inputs for the score. The higher the score, the more potential security exposure would the privilege pose and thus the higher the remediation priority. The resulting risk values allow for ranking of detected excessive privileges and differentiate between critical findings from lower risk access conditions. This risk-based prioritization can help security administrators efficiently review many permissions, and can offer enterprise cloud environments a foundation for the automated classification and controlled privilege remediation.

4. CLOUD PERMISSION AND IDENTITY DATA MODELING




Fig 2 - Cloud Identity-Role-Permission Relationship and Feature Construction Model

4.1. Role and Permission Representation

It represents cloud roles and permissions as structured entities to facilitate systematic analysis of access rights in enterprise environments. Each role has its own permissions, operations, resource scope and own identity types, and each permission is represented by operation, resource and sensitivity level. [7] This representation allows the framework to differentiate between low-risk and high-risk permissions, and recognizes wide or over-assigned access. There are further contextual attributes that are linked to permissions including administrative capability, criticality of resource, frequency of usage and scope of access. The framework offers a consistent framework for analyzing access configurations and identifying privileges in excess of what an identity needs for its operational requirements.

4.2. User Role-Permission Relationships

The user-role-permission relationship model is used to represent the way that identities gain access to cloud resources via directly assigned roles and/or inherited roles and associated permissions. [8] One or more roles are assigned to each user, application, or service account, and

operations and resources are assigned to each role. The framework evaluates them and detects excessive privilege conditions, such as unnecessary role assignments, overlapping permissions, inherited access and permissions that do not match with observed activities. These relationships can be modeled as connected entities, so that the access dependencies are addressed instead of access permissions. This relationship-based representation then facilitates better identification of unwanted privileges in complex enterprise cloud IAM scenarios.

4.3. Feature Construction

Feature construction converts raw identity, permission, resource, role and access-event data into meaningful attributes for excessive-privilege detection and risk analysis. The following are constructed features that include permission usage frequency, access recency, number of assigned permissions, role breadth, resource sensitivity, privilege scope, unused-permission ratio, administrative capability, identity type, and deviations between assigned and observed access. There can be other relationship-based attributes that indicate how many roles an identity has, how much permission overlap exists between roles, and how much access to critical resources an identity has. These features allow for a multi-dimensional view of privilege behavior and configuration, which helps the proposed framework differentiate between the legitimate use of privilege and potentially overprivileged instances, and facilitates later risk scoring and automated classification of instances.

5. AUTOMATED EXCESSIVE PRIVILEGE DETECTION METHODOLOGY




Fig 3 - Automated Excessive Privilege Detection and Risk-Scoring Workflow

5.1. Privilege Detection Pipeline

The proposed privilege detection pipeline offers a systematic process on how to detect excessive privileges in enterprise cloud environments. [9] First, identity, role, permission, resource and access-event data is gathered from the cloud infrastructure and is mapped to a single representation. The framework then examines how permissions are set against how they are being used, contextual factors like resource sensitivity and privilege scope, and then investigates permissions that could be higher than what might be considered a legitimate operational requirement. These are then analyzed against a risk scoring system and categorized by severity of

candidate excess privileges. The final output includes prioritized findings of privileges that can be forwarded to security administrators or to automated remediation processes for access review and/or reduction.

5.2. Permission-to-Usage Analysis

Permission-to-usage analysis consists in comparing the access operations that are actually executed in a specified period of observation with the permissions that are granted to each identity. Each user, application or service account is classified as frequently used, occasionally used, rarely used or not used. [10] The analysis also takes into account the connection between the granted permission and the accessed resource, to decide if the granted permission is proportional to its operational use. The differences in assigned permissions and observed permissions are not necessarily treated as permissions used inappropriately, because there are some permissions that may be necessary for rare but valid activities. This will help the framework determine when access is unnecessary and limit the false flagging of access due to variability in cloud workloads.

5.3. Context-Aware Privilege Analysis

Context-aware privilege analysis considers access rights based on the operational context and security context of each identity and permission. The framework takes into account user or service-account type, assigned role, resource sensitivity, scope of the permissions, access frequency, access time, access location (if available), and historical usage behavior. [11] A permission given to an administrative identity may thus be assessed differently to the same permission granted to a normal application account. This contextual analysis enhances the accuracy of detecting privilege access identifying whether a privilege is being used to access a specific resource in a way that is authorized or not that is necessary to support the reliable detection of excessive privileges in a dynamic cloud environment.

5.4. Privilege Risk Score Calculation

The privilege risk score is a quantitative measure of the security risk that can be posed by a detected excessive privilege. The framework normalizes concepts such as permission sensitivity, resource criticality, privilege scope, usage frequency, access recency, identity characteristics and deviation from expected behaviour into a single risk value. A generalized risk formulation is: $(R_p = \sum_{i=1}^n w_i f_i)$, where R_p is the privilege risk score, f_i is a single risk feature and w_i is the weight of the risk feature. Higher scores mean more potential exposure, and higher remediation priority. This risk-based mechanism allows a consistent decision making process in the privilege review and remediation process and systematically ranks large sets of detected privileges.

5.5. Detection Algorithm

The proposed detection algorithm involves a series of data preparation, permission analysis, contextual evaluation and risk assessment processes to process identity and access information of cloud. [12] The algorithm for each identity captures assigned roles/permissions, assigns them to resources that are accessible, and compares the privileges obtained to the access events observed.

Unused, infrequently used, overly specific or sensitive, or mismatched permissions are flagged as potential excess permissions. The algorithm then assigns a risk score to each candidate, and is able to classify the candidate's risk level as either preprogrammed or learned, according to certain decision criteria. The prioritized list gives an automated foundation for privilege review, ongoing monitoring and controlled remediation, while reducing reliance on manual privilege auditing.

6. MACHINE LEARNING-BASED PRIVILEGE RISK ANALYSIS




Fig 4 - Machine-Learning-Based Privilege Risk Analysis Framework

6.1. Machine Learning Problem Formulation

The excessive privilege detection task is redefined as a supervised machine-learning classification problem where each identity – permission relationship is analyzed and classified based on the privilege status. [13] The learning model takes a feature vector describing "who" has used which role, access to which resources, sensitivity to which resources, and "how" the resources are accessed to determine if the access is legitimate or if it is excessive. The model is represented as $(y=f(X))$, where (X) refers to the extracted privilege related features, (f) is the trained classification function, and (y) is the predicted privilege category. This formulation allows the proposed framework to learn relationship between historical access characteristics and privilege outcomes, which can help to facilitate automatic identification of excessive privileges in dynamic enterprise cloud environments.

6.2. Input Feature Selection

Input feature selection: It helps to reduce the information that will be computed or used by the system to find the most relevant features to know to distinguish between legitimate and excessive privileges. The suggested framework takes into account characteristics based on identity attributes, role assignments, permission scope, access frequency, access recency, resource sensitivity,

administrative capabilities, and deviations between assigned and observed access. Using statistical analysis or model-based feature-selection techniques, features with limited predictive value or with too much correlation can be eliminated. Choosing informative privilege-related features can increase detection accuracy, since the classifier can be trained to pay attention to features that contribute to meaningful evidence of excess privilege; and by improving efficiency of model learning.

6.3. Permission Usage Features

Permission usage characteristics describe what uses an identity has made of a permission recently and often. The following features are important: access frequency, last-access time, number of resources accessed, operation diversity, unused-permission indicators, and ratio of assigned permissions to those used. [14] The features serve as behavioural indicators of whether a permission has been granted to represent a legitimate use of the system. For instance, a permission that is enabled for a long duration and never used may be a good candidate for excessive-privilege analysis, whereas a permission that is used continually could have a sound operational basis. Adding the capability to access privileges with usage features allows the machine-learning model to analyze privileges based on both configuration and observed access.

6.4. Privilege Risk Prediction

The likelihood and severity of excessive access for each identity–permission relationship is estimated through privilege risk prediction, which leverages the trained machine-learning model. The prediction process is fed with the feature vector that is created and outputs a privilege classification, a probability value or both, that reflects the likelihood of the class being assigned in the permission. [15] The prediction covers various aspects of access risk, including permission use, identity behavior, role attributes, privilege scope and resource sensitivity. The risk value could then be added to the rule-based risk factors defined in the detection methodology to prioritize privileges for review. This helps the framework go beyond mere binary detection and help to identify excessive privileges in large-scale cloud settings with a risk awareness.

6.5. Model Training

Model training consists of the creation of a machine-learning classifier by feeding it labeled privilege records, which include identity, permission, role, resource, and access-behavior features. The data available is pre-processed for missing data, normalized with numbers as needed, encoded as categorical variables, and potential class imbalance between legitimate and excessive privileges is handled. The training process is a learning process that minimizes a suitable learning objective in order to learn the relationship between the features and their classifications related to privileges. Several different classifiers including decision-tree based, ensemble methods, or other appropriate classifiers may be assessed to find an effective model for the framework proposed. The trained model is then tested on unseen privilege records to assess the generalization ability.

6.6. Explainability of Privilege Predictions

The proposed framework includes explainability for security administrators to give them understandable reasons for each excessive-privilege prediction. [16] Rather than simply a classification or numerical risk score, the system identifies the characteristics that most strongly influence the predicted result, such as unusually wide permission scope, infrequent usage, access to sensitive resources, or any unusual behavior pattern. Some explanation techniques such as feature importance and local explanation can be used to understand the individual prediction and what made a given identity-permission pair a high-risk prediction. This transparency facilitates human verification, enhances trust in automated privilege analysis, and allows administrators to determine what to do before taking any action to reduce privileges or remediate them.

7. EXPERIMENTAL SETUP AND EVALUATION

7.1. Cloud Dataset Description

The experimental data is representative enterprise cloud identity and access data including user data, service account data, role data, permissions, resource data and access activity data. There are attributes with privileges for each record: identity type, assigned role, scope of permission, sensitivity of resources, frequency of access, recentness of access, and observed operations. There are both valid and over-privileged privilege examples in the dataset for supervised detection and performance evaluation. [17] If real enterprise datasets are unavailable due to security and privacy concerns, a controlled synthetic dataset can be created to simulate realistic configurations and access patterns of cloud IAM. The dataset is constructed to have a wide range of identities, permissions, roles, resources, and usage, which can be used to test the scalability and effectiveness of the proposed framework.

7.2. Dataset Preprocessing

The collected cloud IAM data is processed into a standard format for privilege analysis and machine-learning experiments. This is achieved by data cleansing for duplicate data, imputing missing or inconsistent data, standardizing numerical attributes, encoding categorical identity and permission characteristics, and converting access-event data into usage features. These matching records are then converted into identity-role-permission relationships, and the records are augmented with other contextual attributes like resource sensitivity and privilege scope. If necessary, the dataset is subdivided into training, validation and testing groups with a suitable distribution of instances of legitimate and excessive privilege. These pre-processing steps enhance data quality, suppress noise and achieve that the experimental results are a good representation of the proposed detection methodology.

7.3. Baseline Methods

Several baseline approaches are used for evaluation purposes to test the advantages of automated contextual privilege analysis over conventional approaches. Static rule-based detection, unused-permission analysis, role-based privilege analysis and conventional machine-learning classification methods can be included as baselines. [18] Wherever possible, the same input

information is used for each baseline, and the same experimental conditions are applied to each baseline. While the static approaches mainly rely on pre-defined access rules to identify violations, the usage-based approaches target permission activity and the machine-learning approaches rely on baselines which learn privilege patterns from extracted features. When evaluating these methods with the proposed framework, there is evidence of the added value of a combination of permission configuration, access behavior, context information and risk-based analysis.

7.4. Evaluation Metrics

Several metrics of classification and operation are used to test the effectiveness of the proposed excessive privilege detection framework. Precision is the percentage of detected excessive privileges which are truly excessive, and accuracy is the percentage of all privilege instances correctly classified. Recall measures the framework's ability to detect excessive privileges in the dataset, and the F1 score is a compromise measure between precision and recall. Also, false-positive and false-negative rates are examined to study the wrong privilege classification, and detection time and computational overhead are taken into account to evaluate the practical scalability. Multiple measures give a fuller picture than a single measure of a pupil's performance.

7.5. False-Positive and False-Negative Analysis

Analysis of false-positives and false-negatives ensures that the proposed framework can adequately classify legitimate privileges from excess access. A false positive is when a valid permission is wrongly recognised as a bad permission that could lead to unnecessary access reviews or privilege reduction. False Negative: When a privilege is incorrectly identified as valid and a user is granted access when they shouldn't. The analysis focuses on these errors by type of identity, by type of permission, by sensitivity of the resource, and by usage patterns, looking for circumstances where detection performance might suffer. In enterprise cloud scenarios, reducing false-positive and false-negative rates, as well as unnecessary privileges is critical to effective and practical privilege governance.

8. RESULTS AND DISCUSSION

Table 3: Performance Comparison of Privilege Detection Methods

Detection Method	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	False-Positive Rate (%)	False-Negative Rate (%)	Detection Time
Static Rule-Based Detection	84.72	82.15	79.68	80.90	10.84	20.32	185 ms
Unused Permission Analysis	87.36	85.47	83.92	84.69	8.76	16.08	162 ms
Role-Based Privilege Analysis	89.14	87.63	86.21	86.92	7.41	13.79	174 ms

Conventional ML Model	93.28	92.17	91.46	91.81	4.63	8.54	148 ms
Proposed Framework	96.84	96.21	95.73	95.97	2.81	4.27	126 ms




Fig 5 - Graphical Performance Comparison of Privilege Detection Methods

8.1. Overall Detection Performance

The effectiveness of classification, reliability of detection, and efficiency of operation are used to assess the overall detection performance of the proposed Automated Privilege Detection Framework. [19] The accuracy, precision, recall, F1-score, false-positive rate, false-negative rate, and detection time are used to analyze the experimental results and measure the framework's ability to differentiate between legitimate permissions and excessive privileges. By analyzing multiple elements of permission use, identity behavior, role characteristics, sensitivity to resources, and contextual risk information, the framework can detect potentially excessive access in various cloud identity scenarios. The results highlight how much automated privilege analysis can help with continuous least-privilege enforcement, while minimizing reliance on manual access review activities. The overall assessment also serves as a guideline for the suitability of a framework in large scale enterprise cloud environments.

8.2. Comparison With Baseline Approaches.

The proposed framework is evaluated against the baseline approaches using conventional rules, permission usage and machine learning approaches in similar experimental settings. We are comparing the accuracy, precision, recall, F1-score, false positive rate, false negative rate, and the computational requirements to find differences in effectiveness of detection and operational efficiency. Traditional methods can be effective in detecting pre-determined privilege abuses, but can be ineffective in the face of changing user behavior and access needs. The proposed framework, on the other hand, integrates permission configuration, actual usage behavior, contextual characteristics, and risk analysis to allow for a broader understanding of excessive privileges. The comparative results are hence used to show the advantages and disadvantages of the proposed approach over the existing privilege detection techniques.

8.3. ML Model Performance Comparison

The machine-learning part is tested by the comparison of a set of classification models, which have the same feature set with respect to privilege and use the same experimental data set. The accuracy, precision, recall, F1-score, false-positive rate and computational efficiency are used to assess candidate models for their suitability in predicting excessive privileges. The complexity of tree-based and ensemble models can well describe the nonlinear relationship between permission usage, identity characteristics, role assignments, resource sensitivity and behavioral patterns, whereas simpler models can yield lower computational complexity and higher interpretability. The comparison will lead to the identification of a model that offers a good balance of effectiveness in detection, reliability in prediction, explanation of the model, and computational costs. The model is then embedded in the proposed model for automatic privilege risk prediction and continuous analysis.

9. SECURITY ANALYSIS AND PRACTICAL CONSIDERATIONS

9.1. Security Effectiveness

The effectiveness of the proposed Automated Privilege Detection Framework for security relies on its ability to detect unnecessary, excessive, and high-risk permissions, as well as providing continuous enforcement of the least-privilege principle. The framework leverages privilege risk scores, identity characteristics, access behavior and permission configurations, and resource sensitivity to create a multi-dimensional view of cloud access exposure. Earlier detection of excessive privileges can minimize unnecessary attack surfaces and the damage that can be done with a compromised identity. The framework also facilitates ongoing monitoring to ensure that access patterns and cloud environments evolve and privilege risks are identified. The features highlight how automated privilege analysis can help enhance enterprise cloud security and enhance identity governance.

9.2. Scalability in Enterprise Cloud Environments

The proposed framework is evaluated for scalability by increasing number of identities, roles, permissions, resources and access events without a significant drop in detection performance. The modular design enables the scalable data collection, processing, privilege analysis, machine learning prediction and risk assessment, giving the ability to scale individual components based on workloads. In addition, automated feature extraction and a risk-based prioritization can further minimize the time spent on large permission sets. The framework should make sense in a dynamic enterprise cloud environment where identities and access relationships are constantly evolving. Other factors to consider in terms of scalability are the computational burden of privilege analysis using machine learning, data-processing throughput, memory usage, and processing time.

10. CONCLUSION AND FUTURE WORK

10.1. Summary of Findings

The proposed Automated Privilege Detection Framework uses a combination of identity information, role and permission configurations, access behavior, resource sensitivity and machine-

learning based risk analysis to automatically uncover excessive and unnecessary privileges in enterprise cloud environments. The framework analyzes the use and matches up of assigned privileges to actual use, determining if there is excessive access and prioritizes results based on their security risk. The experimental evaluation framework is designed to evaluate the effectiveness, reliability and feasibility to operationally use the classification-based approach to assess the results of the detection. Overall, the proposed approach is a systematic mechanism for continuous privilege analysis and enforces better least privilege principle in dynamic cloud environments.

10.2. Major Contributions

The key achievement of this study is the creation of an automatic system that is able to combine data from cloud IAM, permission-to-usage analysis, contextual privilege evaluation, risk scoring, and the use of machine-learning classification to detect excessive privileges. The structure allows for a clear representation of the identity-role-permission relationships and includes behavioral and resource-based capabilities for better identification of potentially unneeded access. Furthermore, the risk-scoring system allows privileges identified to be prioritized based on their potential impact on security, providing for better and more efficient privilege governance and remediation. The overall framework aims to minimize reliance on manual access audits and to offer a scalable framework for ongoing privilege monitoring in enterprise cloud settings.

10.3. Future Research Directions

The proposed framework can be further extended by integrating real-world enterprise cloud data sets and adaptive machine-learning models with online learning methods that are able to adapt to changing access patterns. Further investigation can focus on explainable artificial intelligence, graph-based privilege modeling, federated learning, and enhanced anomaly-detection methods to enhance the accuracy of detection and ability to generalize across diverse cloud platforms. Seamless integration with zero-trust architectures and automated identity-governance solutions can also support ongoing privilege verification and remediation. Expansion of the framework to multi-cloud and hybrid-cloud environments and computational efficiency, privacy and adversarial robustness, and real-time deployment requirements should also be explored in future studies.

REFERENCES

- [1] Ross, R., Pillitteri, V., Dempsey, K., Riddle, M., & Guissanie, G. (2019). Protecting controlled unclassified information in nonfederal systems and organizations (No. NIST Special Publication (SP) 800-171 Rev. 2 (Draft)). National Institute of Standards and Technology.
- [2] Hu, V. C., Ferraiolo, D., Kuhn, R., Friedman, A. R., Lang, A. J., Cogdell, M. M., ... & Scarfone, K. (2013). Guide to attribute based access control (abac) definition and considerations (draft). NIST special publication, 800(162), 1-54.
- [3] Hu, V. C., Ferraiolo, D., & Kuhn, D. R. (2006). Assessment of access control systems (Vol. 76). Gaithersburg, MD: US Department of Commerce, National Institute of Standards and Technology.
- [4] Ferraiolo, D., Chandramouli, R., Hu, V., & Kuhn, R. (2016). A comparison of attribute based access control (ABAC) standards for data service applications. NIST Special Publication, 800(2016), 178.
- [5] Yu, S., Wang, C., Ren, K., & Lou, W. (2010, March). Achieving secure, scalable, and fine-grained data access control in cloud computing. In 2010 proceedings ieee infocom (pp. 1-9). IEEE.

- [6] Krutz, R. L., & Vines, R. D. (2010). *Cloud security: A comprehensive guide to secure cloud computing*. Wiley Publishing.
- [7] Rose, S., Borchert, O., Mitchell, S., & Connelly, S. (2019). *Zero trust architecture* (NIST Special Publication 800-207, Initial Public Draft). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.SP.800-207-draft>
- [8] Jensen, M., Schwenk, J., Gruschka, N., & Iacono, L. L. (2009, September). On technical security issues in cloud computing. In *2009 IEEE international conference on cloud computing* (pp. 109-116). IEEE.
- [9] Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM computing surveys (CSUR)*, 41(3), 1-58.
- [10] Andress, J. (2014). *The basics of information security: understanding the fundamentals of InfoSec in theory and practice*. Syngress.
- [11] Chandramouli, R., Ferraiolo, D., & Kuhn, D. (2007). *Role-Based Access Control*. Artech.
- [12] Sabelfeld, A., & Myers, A. C. (2003). Language-based information-flow security. *IEEE Journal on selected areas in communications*, 21(1), 5-19.
- [13] Sanders, M., & Yue, C. (2017, October). Automated least privileges in cloud-based web services. In *Proceedings of the fifth ACM/IEEE Workshop on Hot Topics in Web Systems and Technologies* (pp. 1-6).
- [14] Hatamian, M., Momen, N., Fritsch, L., & Rannenber, K. (2019). A multilateral privacy impact analysis method for android apps. In *Annual Privacy Forum* (pp. 87-106). Springer, Cham.
- [15] Man, N. D., & Huh, E. N. (2011, December). A collaborative intrusion detection system framework for cloud computing. In *Proceedings of the International Conference on IT Convergence and Security 2011* (pp. 91-109). Dordrecht: Springer Netherlands.
- [16] Sanders, M. W., & Yue, C. (2018, March). Minimizing privilege assignment errors in cloud services. In *Proceedings of the Eighth ACM Conference on Data and Application Security and Privacy* (pp. 2-12).
- [17] Zhou, L., Varadharajan, V., & Hitchens, M. (2013). Achieving secure role-based access control on encrypted data in cloud storage. *IEEE transactions on information forensics and security*, 8(12), 1947-1960.
- [18] Kuhn, R., Coyne, E., & Weil, T. (2010). Adding attributes to role-based access control.
- [19] Covington, M. J., Fogla, P., Zhan, Z., & Ahamad, M. (2002, December). A context-aware security architecture for emerging applications. In *18th Annual Computer Security Applications Conference, 2002. Proceedings.* (pp. 249-258). IEEE.
- [20] Wang, Y., Zheng, J., Sun, C., & Mukkamala, S. (2013, July). Quantitative security risk assessment of android permissions and applications. In *IFIP Annual Conference on Data and Applications Security and Privacy* (pp. 226-241). Berlin, Heidelberg: Springer Berlin Heidelberg.
- [21] Ou, X., & Singhal, A. (2011). *Quantitative security risk assessment of enterprise networks*. Berlin: Springer.
- [22] Celik, Z. B., McDaniel, P., Izmailov, R., Papernot, N., Sheatsley, R., Alvarez, R., & Swami, A. (2018, May). Detection under privileged information. In *Proceedings of the 2018 on Asia Conference on Computer and Communications Security* (pp. 199-206).