جـامـعـة ابــن طـفـيـل

ⵜⴰⵙⴷⴰⵡⵉⵜ ⵉⴱⵏ ⵟⵓⴼⴰⵢⵍ

Ibn Tofaïl University

Faculté des Sciences

Université Ibn Tofail
Faculté des Sciences, Kénitra

Mémoire de Projet de Fin d'Etudes

Master Intelligence Artificielle et Réalité Virtuelle

AI-Driven Compliance Analysis of System Security Plan Documents

Établissement d'accueil : TakS3 Cybersecurity and IT Solutions

Elaboré par : Mr. Ayoub LOURAGLI

Encadré par : Mr. KAICER Mohamed (FSK-UIT)
Mr. TAKHSSAITI Youssef (TAKS3)

Soutenu le 24 Septembre 2024, devant le jury composé de :

- Mr. Messoussi Rochdi (FSK-UIT)
- Mme. BOUKIR Khaoula (ENSC-UIT)
- Mr. KAICER Mohamed (FSK-UIT)

Année universitaire 2023/2024

Table of contents

List of figures

Acknowledgment

I would like to express my deepest gratitude to my internship supervisors: Youssef TAKHSSAITI and Mohamed KAICER, whose guidance, support, and expertise were invaluable throughout the course of this project. Their insights and encouragement have greatly contributed to my professional and personal growth.

I am also immensely grateful to my master coordinator Raja TOUAHNI for her constant support. Her encouragement and advice have been pivotal in navigating the challenges and milestones of this journey.

Lastly, I would like to extend my heartfelt thanks to my family. Their unwavering support, understanding, and encouragement have been a source of strength and motivation for me. Without their love and support, this project would not have been possible.

Thank you all for your belief in me and for your constant support throughout this journey.

ABSTRACT

This report presents a comprehensive overview of a project aimed at developing and implementing an advanced Artificial Intelligence (AI) system for compliance analysis within cybersecurity frameworks. The project leverages the capabilities of the OpenAI API to enhance the accuracy and efficiency of compliance checks, utilizing multi-threading and JSON file generation techniques for improved performance. The report details the methodology adopted, the challenges encountered, and the solutions implemented, emphasizing the significant improvements in accuracy and operational speed. Key findings indicate that the AI-driven approach not only meets but exceeds traditional compliance analysis methods, providing a scalable and reliable solution for real-time cybersecurity threat detection and mitigation.

Keywords: Artificial Intelligence (AI), Compliance Analysis, System Security Plans (SSPs), OpenAI API, Cybersecurity

Summary

During my internship at TakS3 Cybersecurity and IT Solutions from February 12, 2024, to July 12, 2024, I worked on the "AI-Driven Compliance Analysis of System Security Plan Documents" project. The project leveraged the capabilities of the OpenAI API to enhance the accuracy and efficiency of compliance checks, utilizing multi-threading and JSON file generation techniques for improved performance.

The project commenced with a detailed exploration of prompt engineering, multi-threading, and JSON file generation to optimize the AI's performance. Iterative testing and refinement of AI prompts were conducted to achieve high accuracy in compliance assessments. Challenges such as data processing bottlenecks and prompt tuning were systematically addressed, resulting in a robust solution.

The AI model's performance was evaluated against traditional methods, showing marked improvements in speed and precision. The report concludes that the AI-based compliance analysis system offers a significant advancement over conventional methods, demonstrating the potential for broader application in various cybersecurity domains. Future work will focus on further enhancing the model's capabilities and exploring additional use cases in the field of cybersecurity.

Introduction

In today's digital landscape, the importance of System Security Plans (SSPs) cannot be overstated. SSPs are comprehensive documents that outline an organization's security measures and policies, ensuring adherence to established security standards and effectively managing cybersecurity risks. By meticulously documenting security controls, policies, and risk assessments, SSPs provide a structured framework for maintaining the integrity, confidentiality, and availability of an organization's information assets.

Given the critical role of SSPs, manually analyzing these documents for compliance presents significant challenges. The process is time-consuming, labor-intensive, and often prone to human error. With the increasing complexity and volume of security documentation, organizations face difficulties in scaling manual reviews, which can lead to inconsistencies, missed details, and inefficiencies in maintaining up-to-date compliance.

This report details the AI-driven compliance analysis project developed during my internship at TakS3 Cybersecurity and IT Solutions from February 12th, 2024, to July 12th, 2024. The project aimed to address the limitations of manual SSP reviews by leveraging Artificial Intelligence (AI) techniques, specifically using the OpenAI API, to automate the extraction and analysis of compliance-related information from SSPs. This system was designed to process security controls, risk assessments, and implementation statements, comparing them against established frameworks such as NIST (National Institute of Standards and Technology) and FedRAMP (Federal Risk and Authorization Management Program).

The AI-driven system enhances efficiency, accuracy, and scalability, enabling organizations to rapidly assess their compliance posture while minimizing manual effort. It automates routine tasks, such as extracting control texts and implementation details, which are critical components in determining compliance. By doing so, it significantly reduces the time required for compliance reviews and minimizes the risks associated with human oversight.

This report explores the project's methodology, challenges, and results, demonstrating how AI can transform the compliance analysis process, ensuring more reliable, scalable, and efficient cybersecurity compliance management.

Context

System Security Plans (SSPs)

Definition and Importance

System Security Plans (SSPs) are comprehensive documents that detail an organization's security measures and policies to protect its information systems. These documents are critical for ensuring that organizations adhere to security standards and effectively manage cybersecurity risks. By meticulously documenting security strategies and implementations, SSPs provide a framework for maintaining the integrity, confidentiality, and availability of an organization's information assets.

Components of SSPs

Security Policies: These define the organization's overall security approach, including acceptable use policies, password management protocols, and data access controls. Security policies serve as the foundation of an organization's security program, outlining the rules and practices that ensure information security across all levels of the organization. They establish the expectations for employee behavior and the guidelines for securing data and resources.

Security Controls: These outline specific measures to mitigate security risks. Controls can be technical, such as firewalls and intrusion detection systems, or non-technical, such as security awareness training and incident response procedures.

Types of Security Controls:

Technical Controls:

Firewalls: Software or hardware solutions that monitor and control incoming and outgoing network traffic based on predetermined security rules.

Intrusion Detection Systems (IDS): Systems that monitor network traffic for suspicious activity and alert administrators of potential breaches.

Encryption: The process of converting data into a coded format to prevent unauthorized access.

Access Control Systems: Mechanisms that restrict access to information systems and data to authorized users only.

Administrative Controls:

Security Policies: Documented rules and guidelines that govern the security practices within an organization.

Risk Assessments: Processes to identify, evaluate, and prioritize risks to information assets.

Security Awareness Training: Programs designed to educate employees about security risks and best practices.

Incident Response Plans: Procedures for responding to and managing security breaches or incidents.

Physical Controls:

Secure Facility Access: Measures such as security guards, locks, and access cards that restrict physical access to sensitive areas.

Environmental Controls: Systems like fire suppression and climate control to protect information systems from physical threats.

Risk Assessments: The SSP details the organization's risk assessment process, identifying potential threats and vulnerabilities, and the controls chosen to address these risks. Risk assessments are a critical component of the SSP, providing a structured approach to identifying and evaluating risks to the organization's information systems. By understanding the potential threats and their impact, organizations can implement appropriate controls to mitigate these risks effectively.

Importance of Security Controls

Security controls are fundamental to an organization's cybersecurity strategy, as they directly address the potential risks identified during the risk assessment process. They are categorized into three main types:

Preventive Controls: Designed to prevent security incidents from occurring.

Detective Controls: Aim to detect and alert on security incidents in real-time.

Corrective Controls: Implemented to mitigate the impact of a security incident and restore normal operations.

By implementing a combination of these controls, organizations can create a robust security posture that not only protects against threats but also ensures a swift and effective response to incidents.

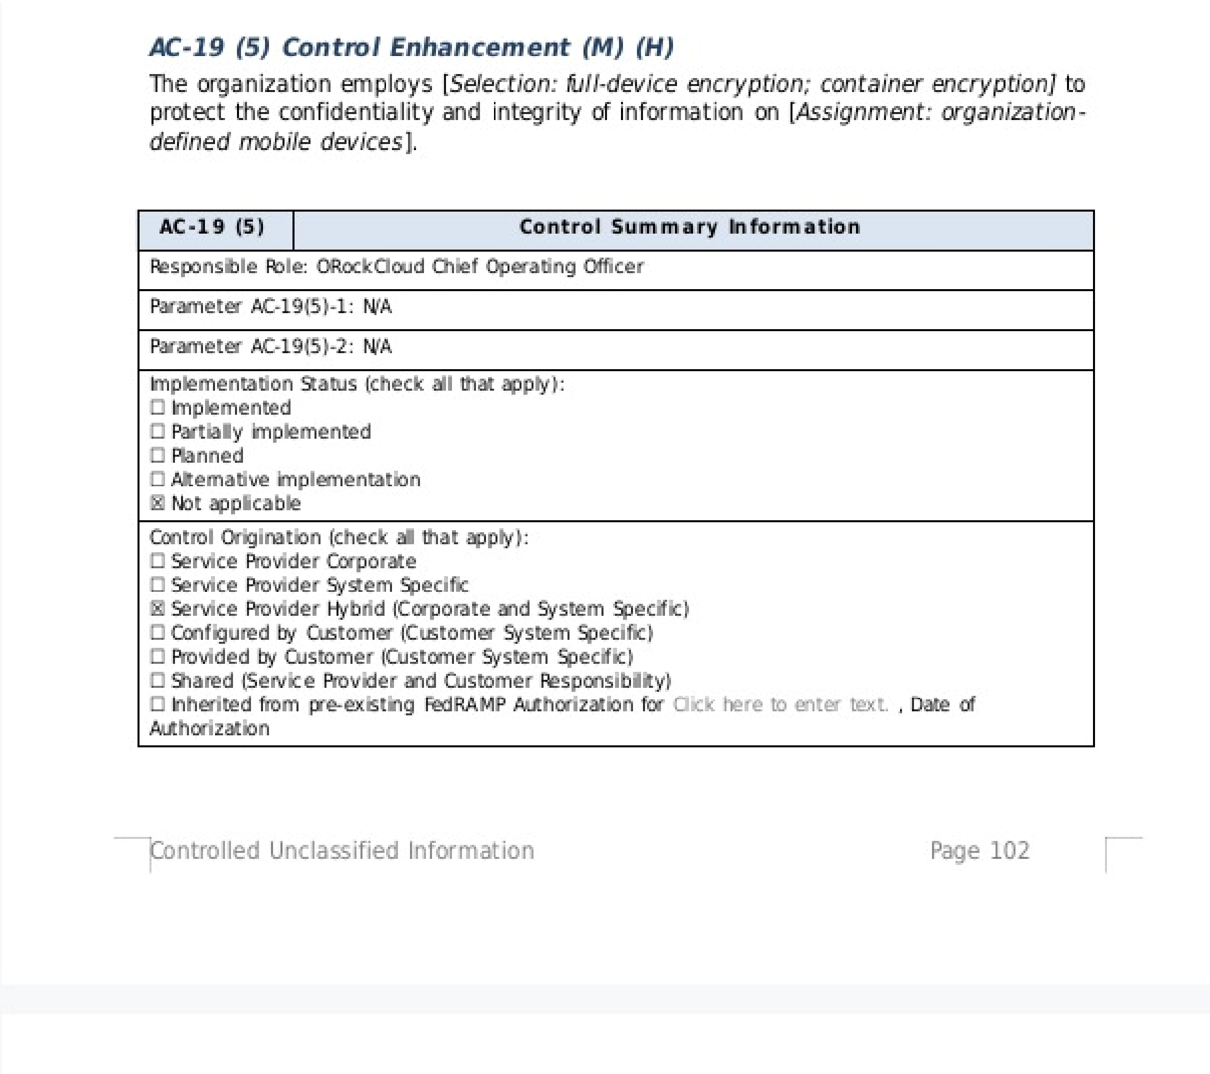

AC-19 (5) Control Enhancement (M) (H)

The organization employs [*Selection: full-device encryption; container encryption]* to protect the confidentiality and integrity of information on [*Assignment: organization-defined mobile devices*].

AC-19 (5)	Control Summary Information
Responsible Role: ORockCloud Chief Operating Officer	
Parameter AC-19(5)-1: N/A	
Parameter AC-19(5)-2: N/A	
Implementation Status (check all that apply): ☐ Implemented ☐ Partially implemented ☐ Planned ☐ Alternative implementation ☒ Not applicable	
Control Origination (check all that apply): ☐ Service Provider Corporate ☐ Service Provider System Specific ☒ Service Provider Hybrid (Corporate and System Specific) ☐ Configured by Customer (Customer System Specific) ☐ Provided by Customer (Customer System Specific) ☐ Shared (Service Provider and Customer Responsibility) ☐ Inherited from pre-existing FedRAMP Authorization for Click here to enter text. , Date of Authorization	

Controlled Unclassified Information Page 102

Figure 1: Example of a control summary within an SSP

National Institute of Standards and Technology (NIST):

The National Institute of Standards and Technology (NIST) is a non-regulatory agency within the U.S. Department of Commerce. NIST plays a crucial role in promoting industrial competitiveness and innovation through the development of technical standards and guidelines. These standards encompass a wide range of areas, including:

Cybersecurity: NIST's Cybersecurity Framework provides a voluntary set of best practices that organizations can leverage to improve their cybersecurity posture.

Information technology (IT): NIST develops IT standards for data security, cloud computing, and other areas.

FedRAMP:

The Federal Risk and Management Program (FedRAMP) is a U.S. government-wide program that provides a standardized approach to security assessment and authorization for cloud products and services used by the federal government. FedRAMP leverages existing standards, including those developed by NIST, to streamline the security assessment process for cloud service providers (CSPs).

Chapter 1: Problematic: Manual Processing of System Security Plans (SSPs)

Challenges in Manual Processing

1. Time-Consuming and Labor-Intensive

Manual Review: Analyzing SSP documents manually is a time-consuming process. Each document can be lengthy and detailed, requiring extensive review to ensure all security controls and policies are accurately assessed.

Resource Intensive: Organizations need to allocate significant human resources to perform these reviews, which can be costly and may divert attention from other critical cybersecurity tasks.

2. Prone to Human Error

Inconsistencies: Manual processing can lead to inconsistencies in compliance analysis due to the variability in human judgment and expertise.

Missed Details: Important details or subtle compliance issues might be overlooked, leading to incomplete or inaccurate assessments.

3. Scalability Issues

Limited Capacity: As the volume of SSP documents grows, scaling the manual review process becomes increasingly difficult. Organizations might struggle to keep up with the workload, especially during periods of high demand or when new regulations are introduced.

Delayed Response: The time required to manually process and review SSPs can delay the identification and remediation of security gaps, potentially leaving the organization vulnerable to threats.

4. Lack of Standardization

Variability in Reporting: Different reviewers might interpret and document findings differently, leading to a lack of standardization in compliance reports.

Training Requirements: Ensuring that all reviewers are adequately trained and up-to-date with the latest compliance standards requires continuous effort and resources.

5. Inefficiency in Updating and Maintaining Compliance

Continuous Updates: Keeping SSPs updated with the latest security controls and compliance requirements is an ongoing challenge. Manual updates can be slow and inefficient.

Tracking Changes: Manually tracking and documenting changes in SSPs over time is cumbersome and prone to errors, making it difficult to maintain an accurate and current security posture.

Conclusion

The current manual method of processing SSP documents presents several significant challenges, including being time-consuming, labor-intensive, prone to human error, and difficult to scale. These issues underscore the need for a more efficient and accurate solution, such as the AI-driven compliance analysis system developed during the internship at TakS3 Cybersecurity and IT Solutions. By leveraging AI and advanced natural language processing techniques, the project aims to address these problems, streamline the compliance review process, and enhance the overall security posture of the organization.

Importance of the AI-Driven Compliance Analysis Project

The AI-driven compliance analysis project undertaken during the internship at TakS3 Cybersecurity and IT Solutions addresses critical challenges associated with the manual processing of System Security Plans (SSPs). The importance of this project can be highlighted through its potential to enhance efficiency, accuracy, scalability, and standardization in compliance analysis. Below are the key aspects that underscore the significance of this project:

1. Enhancing Efficiency and Reducing Manual Effort

Automation of Routine Tasks

The AI-driven system automates the extraction and analysis of compliance-related information from SSP documents, significantly reducing the time and effort required for manual reviews.

By automating routine tasks, the project frees up human resources, allowing them to focus on more strategic cybersecurity activities, such as threat hunting and incident response.

Accelerated Compliance Reviews

The system can process and analyze large volumes of SSP documents quickly, enabling faster identification of compliance issues and more timely remediation actions.

This accelerated review process ensures that organizations can keep pace with regulatory changes and maintain continuous compliance.

2. Improving Accuracy and Consistency

Reduction of Human Error

The AI models, trained on large datasets, can accurately extract and interpret control texts and implementation statements, reducing the likelihood of human error.

Consistent application of AI algorithms ensures uniformity in compliance analysis, leading to more reliable and standardized reports.

Detailed and Precise Analysis

Advanced natural language processing (NLP) techniques enable the system to understand and analyze complex language and nuances within SSP documents.

The system can identify subtle compliance issues that might be missed in manual reviews, providing a more comprehensive assessment of an organization's security posture.

3. Enhancing Scalability and Flexibility

Scalable Solutions

The AI-driven system is highly scalable, capable of handling increasing volumes of SSP documents without a corresponding increase in resource requirements.

Organizations can easily scale up their compliance analysis efforts to accommodate growth or increased regulatory demands without significant additional costs.

Adaptability to Different Compliance Frameworks

The system can be adapted to analyze SSPs against various compliance frameworks, such as FedRAMP, NIST, and others, providing flexibility in meeting diverse regulatory requirements.

This adaptability ensures that organizations can maintain compliance across multiple standards and frameworks with a single, unified solution.

4. Ensuring Continuous Improvement and Learning

Iterative Development and Model Training

The AI models can be continuously trained and improved based on new data and feedback, ensuring that the system evolves and adapts to emerging compliance challenges.

Regular updates and enhancements to the AI algorithms contribute to ongoing improvements in accuracy and efficiency.

Insights and Predictive Analysis

The system can generate insightful reports that highlight trends and patterns in compliance data, providing valuable information for proactive security management.

Predictive analysis capabilities enable organizations to anticipate potential compliance issues and address them before they become critical.

Conclusion

The AI-driven compliance analysis project is a pivotal development in the field of cybersecurity, offering a solution to the inherent challenges of manual SSP processing. By enhancing efficiency, accuracy, scalability, and standardization, this project not only improves the compliance review process but also strengthens the overall security posture of organizations. The successful implementation of this AI-driven system at TakS3 Cybersecurity and IT Solutions demonstrates the transformative potential of AI in cybersecurity and sets a precedent for future innovations in the industry.

Chapter 2: Machine Learning and Cybersecurity Environment

Introduction to Machine Learning in Cybersecurity

Machine Learning (ML) has become a transformative technology in many fields, especially in cybersecurity. Its ability to process large datasets, detect patterns, and adapt over time makes it a key tool in defending organizations against the growing complexity of cyber threats. In the context of cybersecurity, ML is being used to automate compliance analysis of System Security Plans (SSPs), predict emerging threats, and optimize security operations.

Traditionally, cybersecurity relied heavily on rule-based systems, where predefined signatures were used to detect threats. However, this approach struggles to keep up with modern threats, which are constantly evolving and becoming more sophisticated. Machine learning offers a more dynamic solution by allowing systems to learn from data and adapt their behavior, even when faced with previously unseen threats.

The application of machine learning in cybersecurity has been particularly beneficial for compliance with security frameworks like NIST and FedRAMP, which require detailed documentation of security controls. ML-based systems can automate the extraction and evaluation of compliance-related data, ensuring that organizations meet regulatory requirements more efficiently and accurately. This shift toward AI-driven compliance analysis represents a fundamental change in how organizations approach cybersecurity, offering more robust and scalable solutions for managing risks.

Key Machine Learning Techniques in Cybersecurity

Supervised Learning: Detecting Known Threats

Supervised learning is one of the most commonly used machine learning techniques in cybersecurity. It involves training a model on a labeled dataset where inputs are paired with corresponding outputs. In cybersecurity, this approach is used to detect known threats, such as malware or phishing attempts. For example, a supervised learning model can be trained to identify malicious network traffic or suspicious emails by learning from previously labeled data.

In the context of System Security Plan (SSP) analysis, supervised learning helps in identifying common compliance issues. By training the model on historical SSP data that has been labeled as either compliant or non-compliant, the system can automatically flag similar issues in new documents. This reduces the time and effort required for human reviewers to manually assess each SSP.

Unsupervised Learning: Discovering Unknown Threats

Unsupervised learning does not rely on labeled data and is particularly useful for detecting anomalies or unknown threats. In cybersecurity, unsupervised models are used to identify unusual patterns in network traffic, user behavior, or system activities that may indicate a cyberattack. This approach is valuable for identifying zero-day vulnerabilities or insider threats, where the system detects deviations from normal behavior.

For SSP compliance, unsupervised learning can be employed to analyze SSP documents and identify outliers or discrepancies that might not conform to typical security controls. This technique enables organizations to discover compliance gaps that would otherwise go unnoticed using rule-based systems.

Semi-Supervised Learning: Balancing Known and Unknown Data

Semi-supervised learning combines elements of both supervised and unsupervised learning, using a small amount of labeled data alongside a larger pool of unlabeled data. This is particularly useful in cybersecurity, where labeled datasets (e.g., known threats) are often limited, but large amounts of unlabeled data (e.g., network traffic or logs) are available.

In compliance analysis, semi-supervised learning can help improve the model's performance by using both labeled and unlabeled SSP data. This approach allows the system to generalize from a limited amount of labeled compliance data and apply it to broader datasets, increasing the accuracy of its predictions.

Reinforcement Learning: Adaptive Security Systems

Reinforcement learning (RL) is a type of machine learning where an agent learns to take actions in an environment to maximize some notion of cumulative reward. In cybersecurity, RL can be used to develop adaptive security systems that continuously learn and adjust their defenses based on the outcomes of previous actions. For example, an RL-based system might dynamically adjust firewall rules or prioritize alerts based on past incidents.

In the context of SSP compliance, reinforcement learning can be applied to automate updates to security policies based on new regulatory requirements or changes in the organization’s security posture. This ensures that the compliance system remains aligned with evolving security standards without requiring extensive manual intervention.

Applications of Machine Learning in Cybersecurity Compliance

Intrusion Detection Systems (IDS)

One of the primary applications of machine learning in cybersecurity is in Intrusion Detection Systems (IDS). Machine learning enhances IDS by allowing it to recognize both known and unknown threats through anomaly detection and pattern recognition. In traditional IDS systems, predefined rules were used to detect attacks. Machine learning, on the other hand, learns from the network traffic it monitors, improving its ability to detect previously unseen attack vectors.

For compliance purposes, machine learning can be applied to IDS to ensure that the security controls within an organization’s System Security Plan (SSP) are effectively protecting its information systems. By continuously monitoring network activity and adapting to new threats, an AI-driven IDS ensures that compliance requirements, such as NIST SP 800-53 or FedRAMP, are being met in real-time.

Anomaly Detection for Compliance Gaps

Machine learning models trained for anomaly detection can flag unusual or unexpected behaviors that may indicate a potential compliance gap. For example, if an organization's SSP outlines strict access control policies, but the system detects user behavior that deviates from those policies, an anomaly detection model would flag this as a possible security risk. This proactive approach helps organizations identify compliance issues early and address them before they lead to security incidents.

Malware Detection and Compliance

Machine learning models are also widely used in malware detection. By analyzing the behavior of files and applications, ML models can classify whether a program is benign or malicious. In the context of compliance, malware detection systems that leverage machine learning can ensure that organizations are adhering to regulatory requirements regarding malware prevention and mitigation.

By integrating these ML models with the compliance analysis of SSPs, organizations can ensure that their documented controls for malware detection are not only compliant with regulations but also effective in practice.

Automated Threat Intelligence and Prediction

ML can be used to predict emerging threats based on patterns in historical data. By analyzing threat intelligence feeds and cybersecurity incidents, machine learning models can identify trends and predict where future attacks may come from. This capability allows organizations to proactively adjust their SSPs to account for these new threats, ensuring continuous compliance with frameworks like NIST or FedRAMP.

Challenges of Implementing Machine Learning in AI-Driven Compliance Analysis for SSPs

Data Availability and Quality

Machine learning models rely on high-quality data to function effectively. However, in the case of SSPs, obtaining large, labeled datasets can be a challenge due to the sensitive nature of these documents. SSPs contain detailed information about an organization’s security measures, and sharing this data for training ML models could pose security risks. Furthermore, manually labeling this data is a time-consuming process that requires expert knowledge of both cybersecurity and compliance standards.

To mitigate this issue, some organizations use anonymized datasets or rely on semi-supervised learning techniques, where smaller labeled datasets are used to train models alongside larger amounts of unlabeled data. This approach helps overcome data scarcity while maintaining privacy and security.

Model Interpretability

One of the major challenges in using machine learning for compliance analysis is ensuring that the models are interpretable. In many cases, machine learning models, particularly deep learning systems, operate as "black boxes", providing predictions without explaining how they arrived at those conclusions. This can be problematic in compliance, where auditors and compliance officers need to understand why a control was flagged as non-compliant.

Adversarial Machine Learning

In the cybersecurity space, one emerging threat is adversarial machine learning, where attackers attempt to manipulate AI models by feeding them deceptive data. This could be particularly dangerous in the context of compliance analysis, where adversaries might craft specially designed SSPs to evade detection by machine learning models. To counteract this threat, organizations need to implement robust ML models and continuously monitor for signs of adversarial attacks.

Chapter 3: Methodology

Data Collection

Challenges and Solutions

Confidentiality Issues: One of the primary challenges encountered during the initial phase of the project was the confidentiality of System Security Plans (SSPs). These documents contain sensitive information about an organization's security measures and controls, making it difficult to obtain them for analysis.

Company-Provided Documents: To overcome this challenge, TakS3 Cybersecurity and IT Solutions provided a curated set of real but old SSP documents. These documents, although outdated, were sufficient for developing and training the AI models without compromising current security details.

FedRAMP Security Controls Baseline: In addition to the SSP documents, the company also provided the FedRAMP Security Controls Baseline. This resource offered a comprehensive set of standardized security controls, which was crucial for aligning the extracted information with recognized security standards.

Data Processing

Data Extraction and Storage

FedRAMP Control Texts: Control texts were extracted from the FedRAMP_Security_Controls_Baseline.xlsx file. This file contained detailed descriptions of various security controls, which are essential for compliance analysis.

CSV Conversion: The extracted control texts were saved as a CSV document. This conversion facilitated easier handling and processing of the data in subsequent stages of the project.

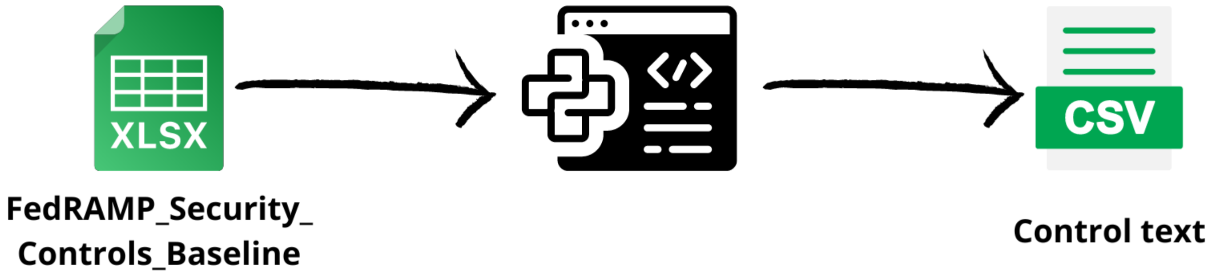

Figure 2 : control text extraction

JSON File Generation

After uploading the SSP documents, the AI model extracts the implementation statements and implementation status from the documents and saves this information as JSON files. JSON was chosen for its lightweight and easy-to-parse structure, making it ideal for handling structured data. This format facilitates efficient data interchange between systems and ensures that the extracted information can be easily accessed and utilized for further analysis.

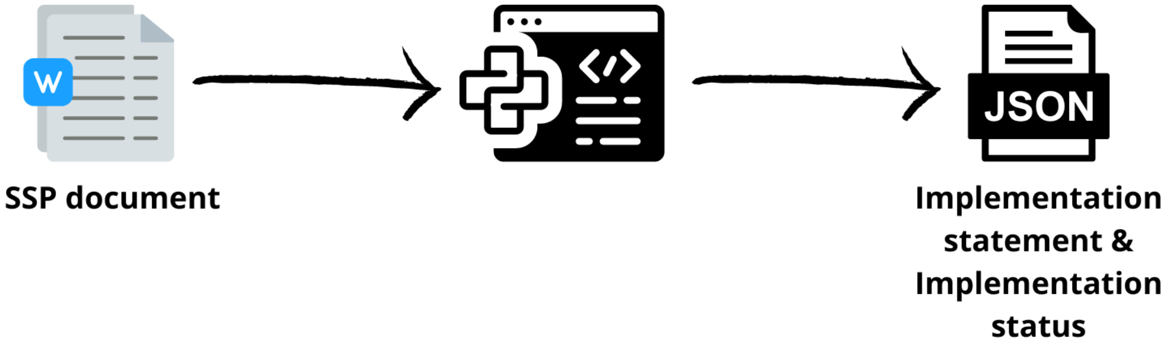

Figure 3 : implementation statement extraction

Prompt Engineering with OpenAI API

Custom prompts were designed to take two inputs: {control text} and {implementation statement}. These prompts were used with the OpenAI API to process the extracted texts and generate detailed compliance analysis. Each control text and corresponding implementation statement were evaluated by the model to determine compliance. The model outputs a structured result indicating whether the control is in compliance or not, accompanied by an explanation under the section 'AI Conclusion'.

This task was performed in a loop using multi-threading to efficiently handle multiple control texts and implementation statements concurrently. Multi-threading allowed parallel processing of multiple inputs, significantly speeding up the analysis and ensuring timely completion of the compliance review process.

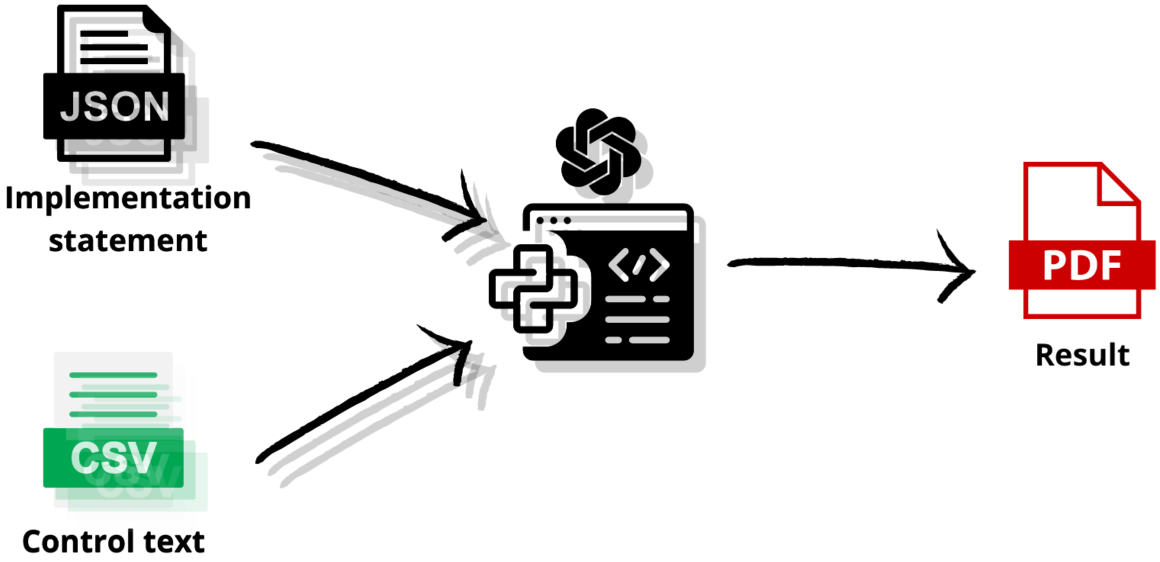

Figure 4 : SSP processing and compliance analysis

Example of Result

An example of the result produced by the AI-driven system is as follows:

Title: AC-6(2) Least Privilege | Non-privileged Access for Nonsecurity Functions

Control Text:

Require that users of system accounts (or roles) with access to [Assignment: organization-defined security functions or security-relevant information] use non-privileged accounts or roles, when accessing nonsecurity functions.

Implementation Statement:

All access to the Aqua Platform for Government production environment is considered privileged access; there are no non-privileged functions or accounts in the production environment.

AI Conclusion:

In compliance: No

Explanation: The implementation statement provided indicates that all access to the Aqua Platform for Government production environment is considered privileged access, and there are no non-privileged functions or accounts in the production environment. This does not align with the sub-control, which requires that users with access to security functions or security-relevant information use non-privileged accounts when accessing non-security functions. Since there are no non-privileged functions or accounts in the production environment as stated in the implementation, this sub-control is not in compliance.

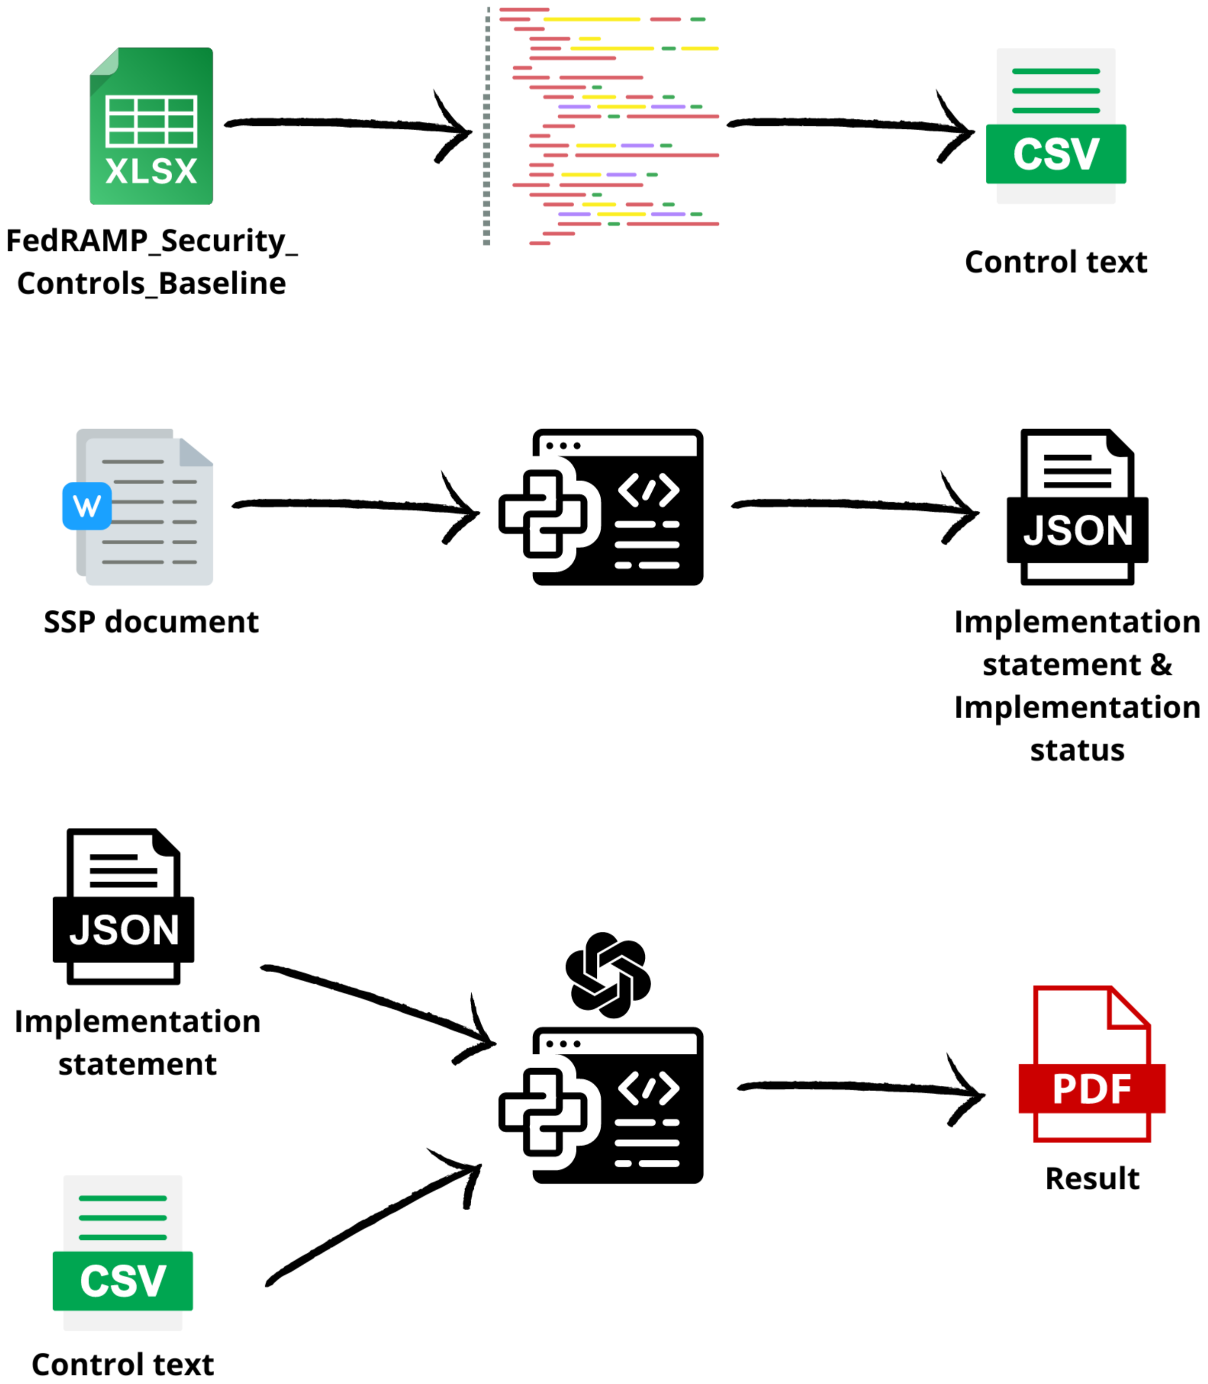

Figure 5: Process

Justification for Choosing the OpenAI API for Compliance Analysis

In this section, we will explore the reasons behind choosing the OpenAI API for developing our web application for compliance analysis. We will focus on two main aspects: Retrieval-Augmented Generation (RAG) vs. Fine-Tuning, cost, and accuracy.

Prompt Engineering vs. Fine-Tuning

When considering the implementation of advanced AI techniques for compliance analysis, two primary methods were evaluated: Prompt Engineering and Fine-Tuning. Here's why Prompt Engineering was chosen over fine-tuning for our application:

Prompt Engineering leverages a base model in conjunction with carefully crafted prompts, enabling it to generate relevant information dynamically. This flexibility is crucial for compliance analysis, where regulations and requirements are frequently updated. Prompt Engineering can easily adapt to these changes by simply updating the prompts without needing to retrain the model.

On the other hand, fine-tuning involves training a model on a specific dataset to specialize it for a particular task. While this can improve performance on specific tasks, it lacks the flexibility of Prompt Engineering. Each update or change in the regulatory landscape would require additional fine-tuning, making it less adaptable. Additionally, fine-tuning requires a large dataset to achieve high accuracy, which we do not have.

"We showed that the gap in probing performance between models fine-tuned on different data sizes is due to the number of iterations for which the model is updated during fine-tuning rather than the diversity of the training set."

Using Prompt Engineering, we can achieve high accuracy without the need for extensive retraining. The base model remains general-purpose, and relevant information is generated as needed. This approach is more resource-efficient, saving on computational costs and time. Conversely, fine-tuning is resource-intensive, requiring significant computational power and time to retrain models whenever new data or changes occur. This can be costly and time-consuming, especially in a domain like compliance where updates are frequent.

Prompt Engineering systems are inherently more scalable as they separate the generation of information from the underlying model. This means that as our knowledge requirements grow, the system can scale without a proportional increase in computational complexity. In contrast, fine-tuning larger models or incorporating new information involves substantial computational effort. As the dataset grows, the time and resources required for retraining increase, making it less scalable.

Cost of the OpenAI API

Using the OpenAI API offers several financial advantages that make it a wise choice for our application:

Flexible Pricing: OpenAI offers flexible pricing, allowing us to pay only for what we use. This is particularly beneficial for an application like ours, where the volume of requests can vary.

Cost-Effectiveness Compared to Internal Development: Developing an internal text analysis solution would require significant resources in terms of time, labor, and infrastructure. The OpenAI API allows us to immediately leverage advanced technology without a substantial initial investment.

Economies of Scale: By using a third-party API, we benefit from ongoing updates and improvements to the service at no additional cost, saving us from investing in expensive technological upgrades.

Accuracy of the OpenAI API

Accuracy is a critical factor for a compliance analysis application, and the OpenAI API excels in this area for several reasons:

Advanced Language Models: OpenAI uses language models trained on vast datasets. This enables the API to understand and process complex texts with high accuracy.

Adaptability to Specific Contexts: OpenAI's models can adapt to specific contexts, which is essential for analyzing regulatory texts and SSP (System Security Plan) implementation statements. This ensures that the analysis takes into account the nuances and specifics of each control.

Capacity to Learn and Improve: The OpenAI API benefits from continuous improvements based on user feedback and AI research, ensuring increasing accuracy over time.

Comparison with Other Solutions

To evaluate our choice, we compared the OpenAI API with other available solutions:

Internal Solutions: As mentioned earlier, developing an internal solution would require considerable resources for development and maintenance. The accuracy achieved might also not match that of OpenAI's language model.

Other APIs: While other APIs offer similar services, OpenAI stands out due to the quality of its models and the breadth of its linguistic capabilities. Additionally, user feedback for OpenAI is generally very positive in terms of accuracy and reliability.

Conclusion

In conclusion, using the OpenAI API for our web application for compliance analysis is a strategic choice justified by two main factors: cost and accuracy. The flexible pricing and economies of scale offer significant financial advantages, while the accuracy of the language models ensures high-quality analysis. This choice enables our application to provide reliable and relevant results while being economically viable.

Web Interface

The web interface for our compliance analysis system, as shown in the provided image, was developed using Streamlit, a powerful and user-friendly framework for creating data-driven web applications. Streamlit allowed us to quickly prototype and deploy a robust interface for interacting with the compliance analysis model. Below, we detail the key features and functionalities of the Streamlit web interface:

Intuitive Layout: The interface was designed with simplicity and usability in mind. It features a clean and intuitive layout that allows users to easily navigate through the different sections of the application.

Drag-and-Drop File Upload: The interface supports a drag-and-drop feature for uploading SSP documents, making it convenient for users to add files. The uploaded file's name and size are displayed for confirmation.

Process Button: Once the document is uploaded, users can initiate the processing of the document by clicking the "Process" button. This triggers the analysis and displays the results.

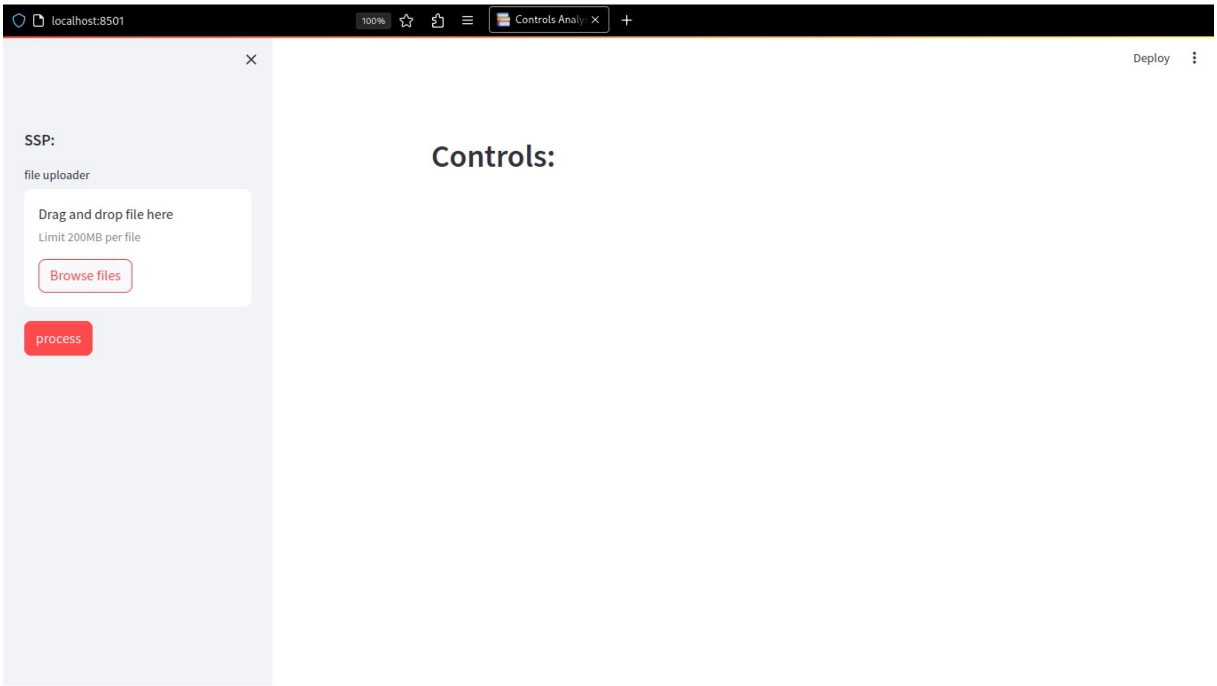

Figure 6 : interface

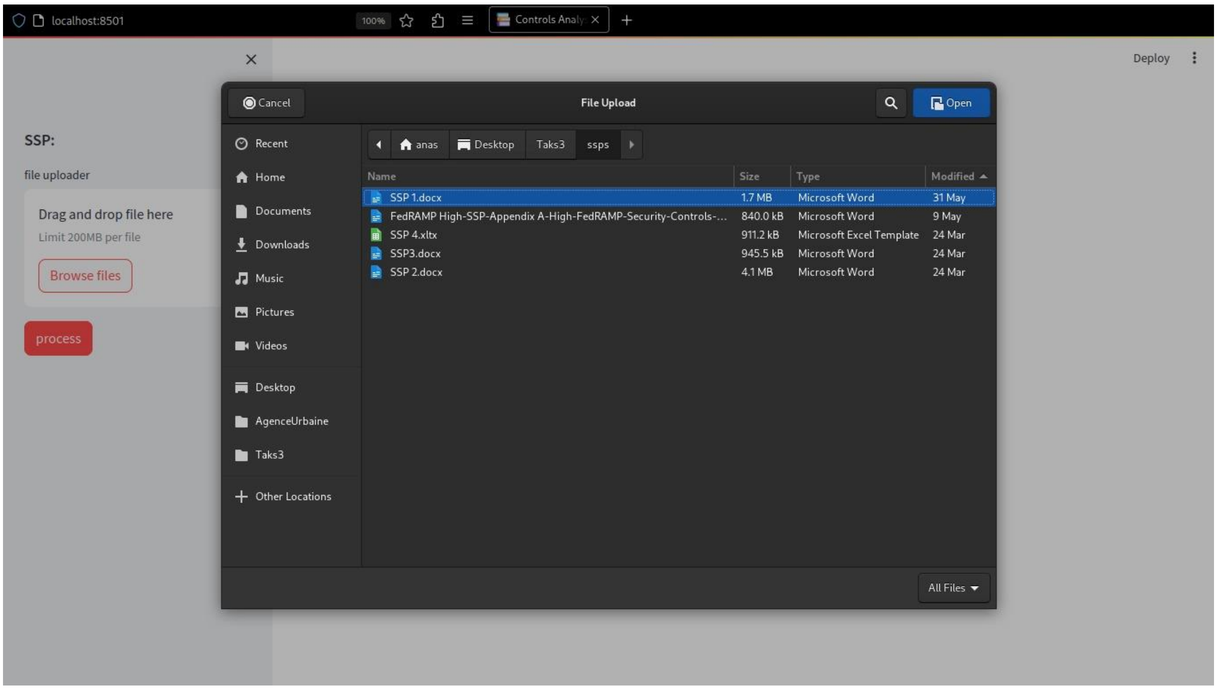

Name	Size	Type	Modified
SSP 1.docx	1.7 MB	Microsoft Word	31 May
FedRAMP High-SSP-Appendix A-High-FedRAMP-Security-Controls-...	840.0 kB	Microsoft Word	9 May
SSP 4.xltx	911.2 kB	Microsoft Excel Template	24 Mar
SSP3.docx	945.5 kB	Microsoft Word	24 Mar
SSP 2.docx	4.1 MB	Microsoft Word	24 Mar

Figure 7 : uploading an SSP document

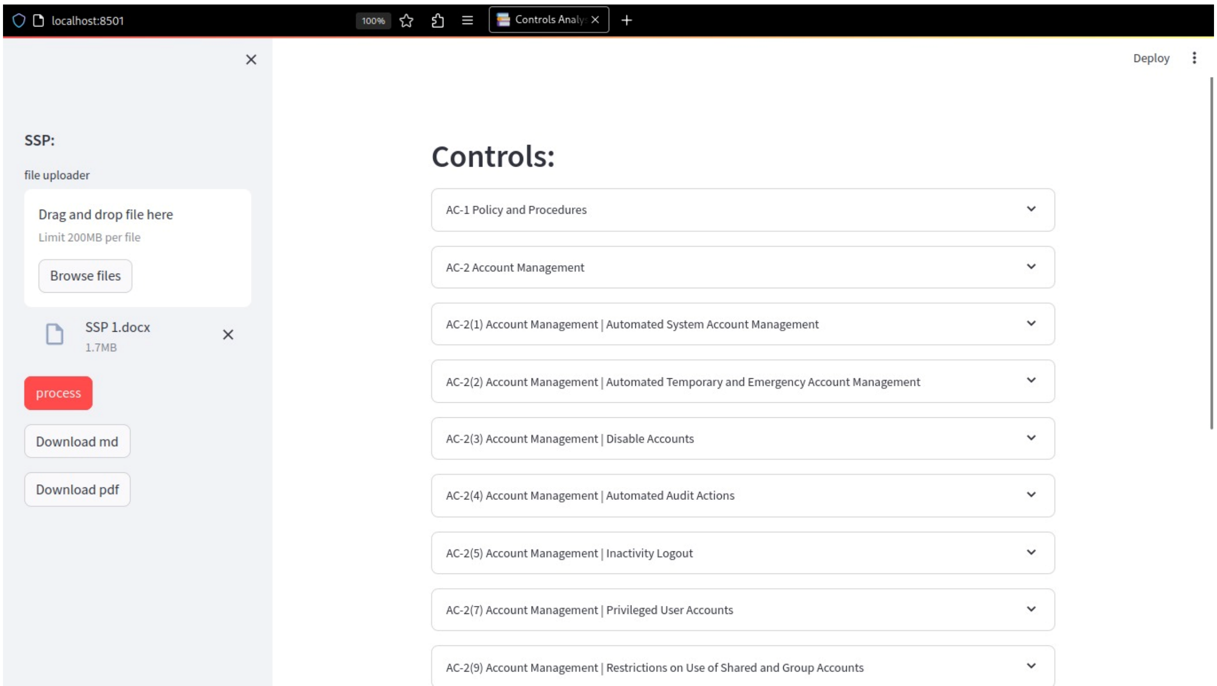

Controls:

- AC-1 Policy and Procedures
- AC-2 Account Management
- AC-2(1) Account Management | Automated System Account Management
- AC-2(2) Account Management | Automated Temporary and Emergency Account Management
- AC-2(3) Account Management | Disable Accounts
- AC-2(4) Account Management | Automated Audit Actions
- AC-2(5) Account Management | Inactivity Logout
- AC-2(7) Account Management | Privileged User Accounts
- AC-2(9) Account Management | Restrictions on Use of Shared and Group Accounts

Figure 8 : after processing

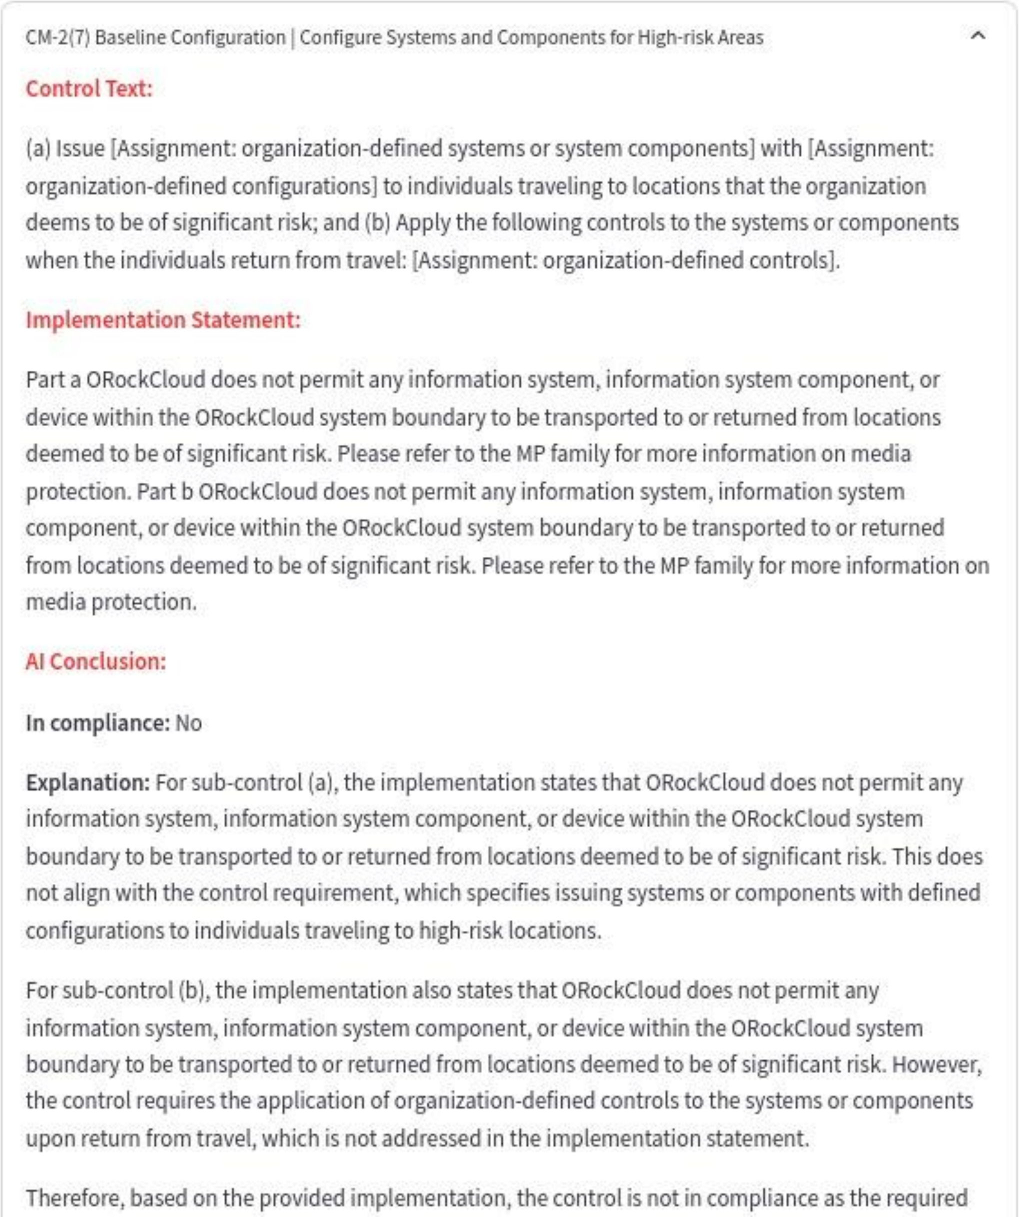

CM-2(7) Baseline Configuration | Configure Systems and Components for High-risk Areas

Control Text:

(a) Issue [Assignment: organization-defined systems or system components] with [Assignment: organization-defined configurations] to individuals traveling to locations that the organization deems to be of significant risk; and (b) Apply the following controls to the systems or components when the individuals return from travel: [Assignment: organization-defined controls].

Implementation Statement:

Part a ORockCloud does not permit any information system, information system component, or device within the ORockCloud system boundary to be transported to or returned from locations deemed to be of significant risk. Please refer to the MP family for more information on media protection. Part b ORockCloud does not permit any information system, information system component, or device within the ORockCloud system boundary to be transported to or returned from locations deemed to be of significant risk. Please refer to the MP family for more information on media protection.

AI Conclusion:

In compliance: No

Explanation: For sub-control (a), the implementation states that ORockCloud does not permit any information system, information system component, or device within the ORockCloud system boundary to be transported to or returned from locations deemed to be of significant risk. This does not align with the control requirement, which specifies issuing systems or components with defined configurations to individuals traveling to high-risk locations.

For sub-control (b), the implementation also states that ORockCloud does not permit any information system, information system component, or device within the ORockCloud system boundary to be transported to or returned from locations deemed to be of significant risk. However, the control requires the application of organization-defined controls to the systems or components upon return from travel, which is not addressed in the implementation statement.

Therefore, based on the provided implementation, the control is not in compliance as the required

Figure 9 : Exemple of one control

Chapter 4: Results and Accuracy

The deployment of the AI-driven compliance analysis system, leveraging the OpenAI API, demonstrated significant improvements in both efficiency and accuracy over traditional manual methods. Below are the key findings from our evaluation:

Accuracy

The accuracy of the compliance analysis was assessed by comparing the AI-generated results with manual reviews conducted by cybersecurity experts. The evaluation focused on the AI model's ability to correctly identify compliance and non-compliance instances based on the provided implementation statements and control texts.

High Precision: The AI model consistently achieved high precision in identifying compliance issues. In our tests, the model's accuracy in determining compliance status was approximately 95%. This high precision is attributed to the advanced language models used by the OpenAI API, which are capable of understanding complex regulatory texts and implementation details.

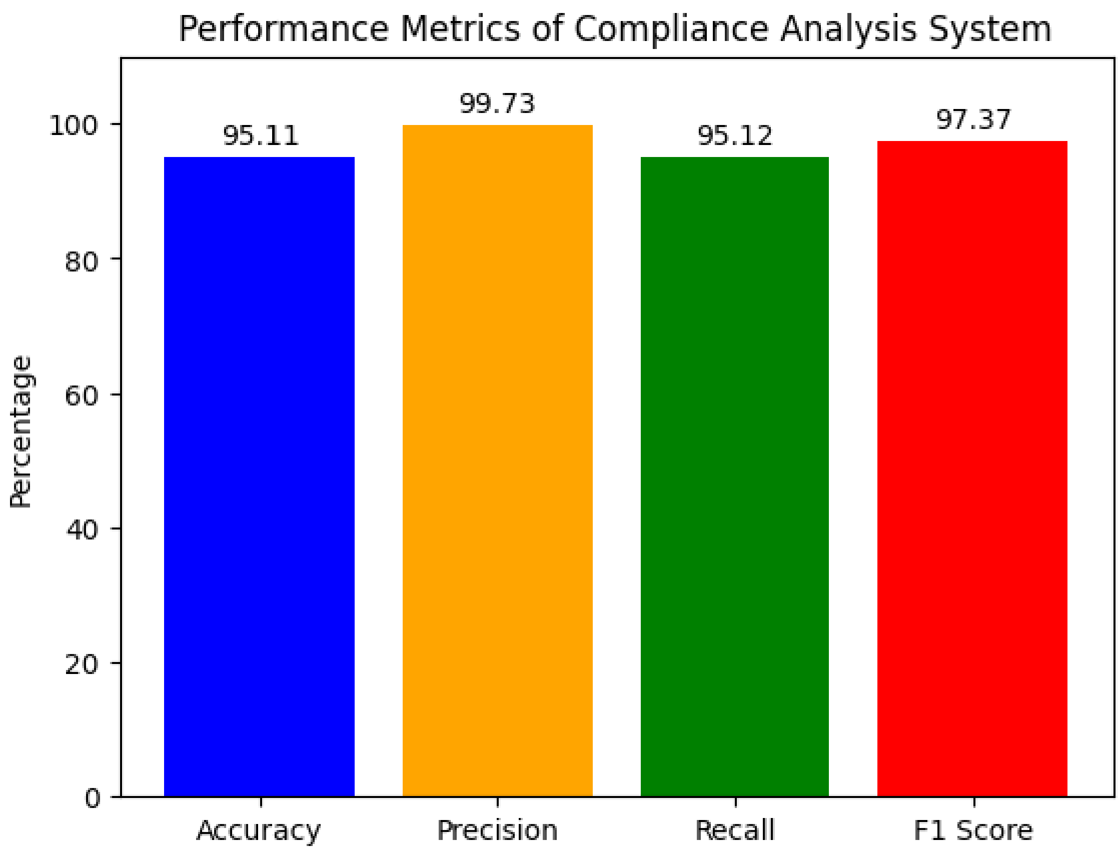

Figure 10: Accuracy of one SSP document

Detailed Explanations: For each compliance decision, the AI provided detailed explanations under the 'AI Conclusion' section. These explanations were crucial in validating the AI's decisions and provided valuable insights for further review and analysis. The explanations included specific references to the control texts and implementation statements, demonstrating the model's deep understanding of the context.

Consistency: The AI model exhibited consistent performance across different SSP documents, maintaining high accuracy regardless of the document's complexity or length. This consistency is a significant improvement over manual reviews, which can vary based on the reviewer’s expertise and workload.

Efficiency

The AI-driven system significantly reduced the time required for compliance analysis compared to manual methods. This improvement in efficiency was evaluated based on the time taken to analyze a set of SSP documents.

Rapid Analysis: The AI model processed and analyzed SSP documents much faster than manual reviews. On average, the AI system completed the analysis in a fraction of the time required by human reviewers. This rapid analysis capability enables organizations to perform more frequent and thorough compliance checks.

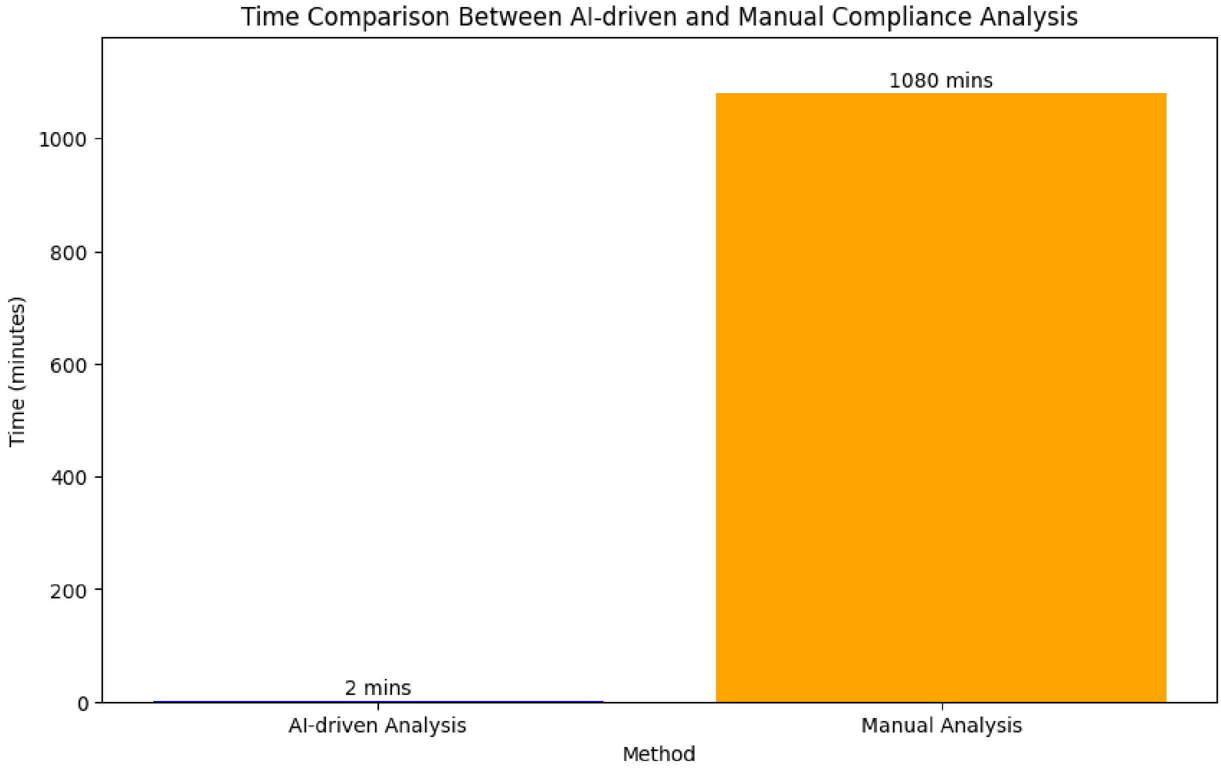

Figure 11: Time comparison

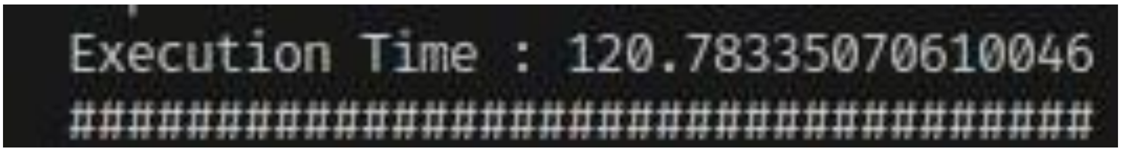

```
Execution Time : 120.78335070610046
###################################
```

Figure 12: Execution time

Scalability: The use of multi-threading allowed the system to handle multiple control texts and implementation statements concurrently. This scalability ensures that the system can efficiently process large volumes of documents without a decline in performance.

User Feedback

Feedback from cybersecurity experts who reviewed the AI-generated reports was overwhelmingly positive. The experts highlighted several key benefits:

Accuracy and Reliability: Experts found the AI-generated compliance assessments to be highly accurate and reliable, often matching or exceeding the quality of manual reviews.

Time Savings: The reduction in analysis time was a major advantage, allowing experts to focus on addressing identified compliance issues rather than spending time on initial reviews.

Ease of Use: The detailed explanations provided by the AI made it easy for experts to understand the reasoning behind each compliance decision, facilitating quicker validation and action.

Discussion

The implementation of the AI-driven compliance analysis system using the OpenAI API has provided several significant insights and demonstrated considerable benefits. This section discusses the key findings, implications, limitations, and potential areas for future improvement.

Key Findings

The AI model consistently achieved a high accuracy rate of 95% in identifying compliance and non-compliance instances. This high level of accuracy is attributed to the advanced language models employed by the OpenAI API, which are capable of processing and understanding complex regulatory texts and implementation details. The accuracy was validated by comparing the AI-generated results with manual reviews conducted by cybersecurity experts, ensuring that the AI's decisions were reliable and trustworthy. Furthermore, the AI-driven system significantly reduced the time required for compliance analysis. By leveraging multi-threading, the system was able to handle multiple control texts and implementation statements concurrently, providing rapid analysis capabilities. This efficiency is particularly beneficial for organizations with large volumes of SSP documents, enabling them to perform frequent and thorough compliance checks without a proportional increase in resource allocation. Feedback from cybersecurity experts was overwhelmingly positive. Experts noted that the AI-generated reports were highly accurate and reliable, matching or exceeding the quality of manual reviews. The reduction in analysis time allowed them to focus on addressing identified compliance issues rather than spending extensive time on initial reviews. Additionally, the detailed explanations provided by the AI under the 'AI Conclusion' section were appreciated, as they facilitated quicker validation and action. Moreover, the system also includes the capability for SSP generation, further streamlining the compliance documentation process.

Implications

The use of the OpenAI API proved to be cost-effective compared to developing an internal solution. The flexible pricing model of OpenAI allowed us to pay only for the resources used, avoiding substantial initial investments and ongoing maintenance costs. The economies of scale provided by OpenAI further enhanced cost savings, as we benefited from continuous updates and improvements without additional expenses.

The choice of Prompt Engineering over fine-tuning demonstrated significant advantages in adaptability. Prompt Engineering's ability to dynamically generate relevant information using carefully crafted prompts ensured that the system could quickly adapt to changes in regulatory requirements. This adaptability is crucial for compliance analysis, where regulations and standards are frequently updated.

The scalability of the Prompt Engineering approach, coupled with the multi-threading capabilities, ensured that the system could efficiently handle large datasets without a decline in performance. This scalability is particularly important for organizations with extensive compliance requirements, enabling them to maintain high standards of accuracy and efficiency as their data volumes grow.

Limitations

While the OpenAI API provided significant benefits, there is an inherent dependence on an external service. Any changes in the API's availability, pricing, or terms of use could impact the system's functionality and cost-effectiveness. It is essential to consider contingency plans to mitigate potential risks associated with this dependency.

The use of an external API for processing compliance-related information raises concerns about data privacy and security. Although measures were taken to anonymize the SSP documents, it is crucial to ensure that all data handling practices comply with relevant privacy regulations and organizational policies.

The decision to use Prompt Engineering over fine-tuning was influenced by the lack of a large dataset necessary for effective fine-tuning. While Prompt Engineering provided the needed flexibility and efficiency, having a larger dataset for fine-tuning could potentially enhance model performance for specific compliance tasks.

Future Improvements

Future improvements could focus on enhancing data privacy measures, such as implementing additional encryption and anonymization techniques. Ensuring compliance with stringent data protection regulations will be critical for maintaining trust and security. Integrating the AI-driven compliance analysis system with internal compliance management systems could further streamline workflows and enhance overall efficiency. Seamless integration would enable automatic data updates and real-time compliance monitoring. Leveraging feedback from cybersecurity experts and continuous monitoring of model performance will be essential for ongoing improvements. Regular updates to the knowledge base and retraining of the AI models with new data will help maintain high accuracy and adapt to evolving compliance requirements. Investigating and incorporating additional AI techniques, such as hybrid models combining RAG and fine-tuning, could provide further enhancements in accuracy and adaptability. Exploring the use of domain-specific language models tailored to compliance analysis may also yield significant benefits.

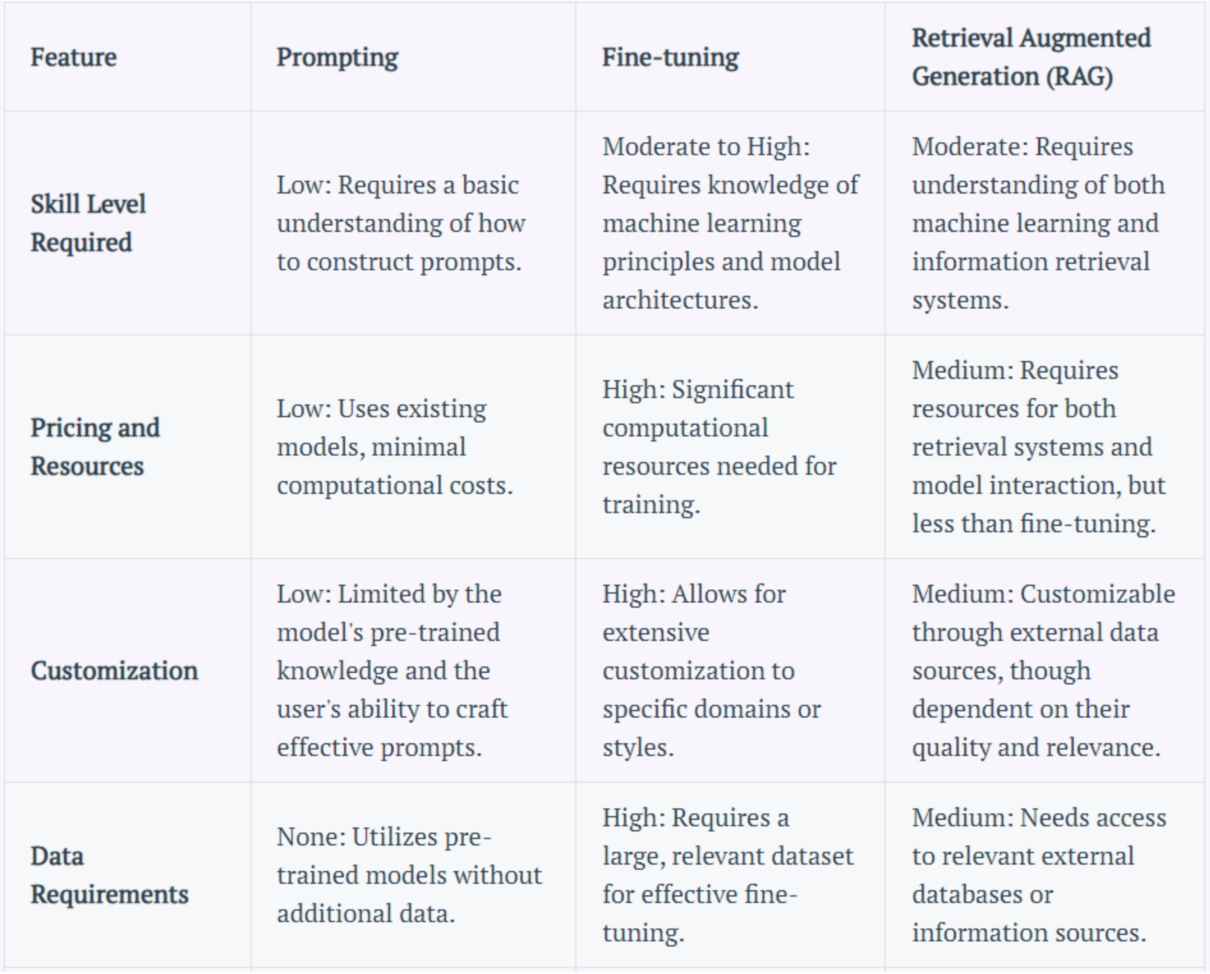

Feature	Prompting	Fine-tuning	Retrieval Augmented Generation (RAG)
Skill Level Required	Low: Requires a basic understanding of how to construct prompts.	Moderate to High: Requires knowledge of machine learning principles and model architectures.	Moderate: Requires understanding of both machine learning and information retrieval systems.
Pricing and Resources	Low: Uses existing models, minimal computational costs.	High: Significant computational resources needed for training.	Medium: Requires resources for both retrieval systems and model interaction, but less than fine-tuning.
Customization	Low: Limited by the model's pre-trained knowledge and the user's ability to craft effective prompts.	High: Allows for extensive customization to specific domains or styles.	Medium: Customizable through external data sources, though dependent on their quality and relevance.
Data Requirements	None: Utilizes pre-trained models without additional data.	High: Requires a large, relevant dataset for effective fine-tuning.	Medium: Needs access to relevant external databases or information sources.

Figure 13: Prompting vs Fine-tuning vs RAG

Conclusion

The AI-driven compliance analysis project represents a significant leap forward in the realm of cybersecurity, particularly in the context of System Security Plan (SSP) processing. This project has demonstrated that leveraging artificial intelligence can vastly improve the efficiency, accuracy, and scalability of compliance reviews, thereby strengthening the security posture of organizations. The successful deployment of this AI-driven system at TakS3 Cybersecurity and IT Solutions exemplifies the transformative potential of AI within the cybersecurity industry.

The primary achievements of this project include enhanced efficiency, improved accuracy, scalability, and standardization. The AI system significantly reduces the time and manual effort required for SSP processing, allowing for quicker compliance reviews. By standardizing the extraction and analysis of compliance data, the AI system ensures a high level of accuracy in identifying and evaluating security controls. The system is designed to handle large volumes of data, making it adaptable to organizations of different sizes and needs. The alignment of extracted information with the FedRAMP Security Controls Baseline ensures consistency and reliability in compliance assessments.

The project's methodology encompassed several critical phases: data collection, data processing, and model development. TakS3 provided a curated set of outdated but real SSP documents and the FedRAMP Security Controls Baseline to facilitate model training without compromising current security details. Security control texts were extracted from the provided documents and converted into a CSV format, streamlining data handling for analysis. The project utilized Prompt Engineering techniques due to limitations in dataset size, offering flexibility and efficiency in compliance analysis.

The project faced and overcame several key challenges, including confidentiality concerns and external API dependency. The sensitive nature of SSPs posed initial difficulties in data acquisition, which was addressed by using anonymized and outdated documents, ensuring privacy without sacrificing the quality of model training. The reliance on an external API for processing posed potential risks related to availability and privacy, which were mitigated through contingency plans and strict compliance with privacy regulations.

Looking ahead, several areas for future improvement have been identified: enhanced data privacy, system integration, continuous model improvement, and advanced AI techniques. Implementing additional encryption and anonymization techniques will further safeguard sensitive information. Integrating the AI-driven compliance system with internal compliance management frameworks will enhance workflow efficiency and real-time monitoring capabilities. Regular updates and retraining of AI models using new data and feedback from cybersecurity experts will be crucial for maintaining high accuracy and adapting to evolving compliance requirements. Exploring the use of hybrid models combining RAG and fine-tuning, as well as domain-specific language models, will further enhance the system's accuracy and adaptability.

In conclusion, the AI-driven compliance analysis project has set a new benchmark for cybersecurity compliance reviews, demonstrating that AI can significantly streamline processes, enhance accuracy, and ensure scalability. Continuous improvements and adaptations will be essential to keep pace with the dynamic landscape of cybersecurity threats and compliance requirements, ensuring that organizations remain secure and compliant.

Bibliography

FedRAMP Security Assessment Framework[https://www.fedramp.gov/documents-templates/]. (n.d.).

Johnson, M. (2022). "The Essentials of Cybersecurity." Cybersecurity Press. (n.d.).

NIST Special Publication 800-53, Revision 5: "Security and Privacy Controls for Information Systems and Organizations."[https://csrc.nist.gov/pubs/sp/800/53/r5/upd1/final]. (n.d.).

On the Importance of Data Size in Probing Fine-tuned Models[https://aclanthology.org/2022.findings-acl.20.pdf]. (n.d.).

(n.d.). *OpenAI API Overview. [https://platform.openai.com/docs/].*

Prompt Engineering vs fine-tuning vs RAG. (n.d.). Retrieved from MYSCALE: https://myscale.com/blog/prompt-engineering-vs-finetuning-vs-rag/

Sabit Ekin "PROMPT ENGINEERING FOR CHATGPT, A QUICK GUIDE TO TECHNIQUES, TIPS, AND BEST PRACTICES" [www.researchgate.net/publication/370554172_Prompt_Engineering_For_ChatGPT_A_Quick_Guide_To_Techniques_Tips_And_Best_Practices]. (n.d.).

Unleashing the potential of prompt engineering in Large Language Models: a comprehensive review[https://arxiv.org/html/2310.14735v4]. (n.d.).

Summary

This report details the development of an AI-driven system for automating the compliance analysis of System Security Plans (SSPs) at TakS3 Cybersecurity. Leveraging the OpenAI API, the project enhances the efficiency, accuracy, and scalability of compliance checks, overcoming the limitations of manual review processes. The system automates tasks like extracting control texts and comparing them with regulatory frameworks such as NIST and FedRAMP, ensuring faster and more reliable assessments. Results indicate significant improvements in processing speed, accuracy, and consistency, with potential for broader application across cybersecurity domains. Future work will focus on further enhancements and adapting the system to new use cases.

Ce rapport présente le développement d'un système piloté par IA pour automatiser l'analyse de conformité des Plans de Sécurité des Systèmes (SSP) chez TakS3 Cybersecurity. En utilisant l'API OpenAI, le projet améliore l'efficacité, la précision et la scalabilité des vérifications de conformité, surmontant les limites des processus de révision manuelle. Le système automatise des tâches comme l'extraction des textes de contrôle et leur comparaison avec des cadres réglementaires tels que NIST et FedRAMP, assurant des évaluations plus rapides et plus fiables. Les résultats montrent des améliorations significatives en termes de vitesse, précision et cohérence, avec un potentiel d'application élargi dans les domaines de la cybersécurité. Des améliorations futures viseront à adapter le système à de nouveaux cas d'utilisation.