International Journal of Engineering Fields (IJEF)
Volume 4, Issue 4, October-December, 2025
Available online at: https://journalofengineering.org
This is an open access article under the CC BY-NC license.

Secure Cloud SaaS Intellectual Property Protection through Generative AI and Behavioral Analytics

Kalpesh Rathod

Director of Engineering, *Lecorpio LLC*, Fremont, CA, USA

ABSTRACT

Cloud-based Software-as-a-Service (SaaS) platforms increasingly host high-value intellectual property (IP), including source code, proprietary algorithms, technical documents, AI models, product designs, and confidential business information. Conventional security mechanisms such as role-based access control, authentication, and data-loss prevention primarily rely on predefined permissions and static rules, making them less effective against compromised identities, insider misuse, behavioral anomalies, and AI-mediated IP extraction. This research proposes a Generative Behavioral Analytics Intellectual Property Protection Framework (GBA-IPPF) **that integrates behavioral analytics, generative artificial intelligence, contextual risk scoring, data-protection mechanisms, and adaptive policy enforcement into a unified SaaS security architecture. The framework continuously collects authentication, API, repository, download, sharing, device, and AI-interaction events; establishes behavioral profiles for users and service identities; analyzes IP sensitivity and semantic relationships using GenAI; and dynamically determines the risk associated with each activity. The proposed implementation was evaluated in a controlled cloud-SaaS experimental environment containing** 120,000 activity events and 8,000 sessions**, covering legitimate behavior and multiple IP-threat scenarios.**

Experimental results demonstrate that GBA-IPPF significantly outperforms RBAC+DLP, Behavioral UEBA, and GenAI Security baselines. The proposed framework achieved 98.21% accuracy, 97.86% precision, 97.44% recall, 97.65% F1-score, and 0.9912 AUC**. Attack-specific detection rates ranged from** 97.5% to 99.3%**, including strong performance against prompt-based extraction and AI-assisted IP reconstruction. Furthermore, GBA-IPPF reduced the false-positive rate to** 4.9%**, detection latency to** 163 ms**, and response latency to** 279 ms**. Most importantly, simulated IP leakage events decreased from** 412 with RBAC+DLP to 73**, representing an approximately** 82.3% reduction**. These findings indicate that combining behavioral intelligence with GenAI-based semantic understanding provides stronger protection against both conventional and emerging SaaS IP threats. The research demonstrates that GBA-IPPF can provide a continuous, context-aware, and adaptive security mechanism capable of detecting suspicious behavior, understanding the sensitivity of affected IP, and automatically enforcing appropriate protective actions.**

Keywords: Security, Cloud, SaaS, Intellectual Property, Protection, Generative AI and Behavioral Analytics

INTRODUCTION

The rapid adoption of cloud-based Software-as-a-Service (SaaS) platforms [1] has fundamentally changed how organizations create, store, process, and exchange intellectual property (IP) [2]. Source code, proprietary algorithms, technical specifications, research outputs, business strategies, machine-learning models, and confidential documents are increasingly maintained within distributed cloud environments rather than traditional enterprise data centers. This transformation provides scalability and accessibility but also increases the attack surface associated with identities, APIs, third-party integrations, remote access, and shared cloud resources. Conventional perimeter-oriented security is increasingly inadequate because users, devices, applications, and resources operate across heterogeneous environments. The Zero Trust Architecture (ZTA) proposed by NIST emphasizes continuous verification of users and devices and protection of individual resources rather than implicit trust based on network location. Consequently, SaaS IP protection requires mechanisms capable of continuously assessing identity, context, behavior, and resource sensitivity rather than relying exclusively on static access permissions [3].

Intellectual-property protection in SaaS environments is particularly challenging because legitimate users may already possess authorized access to valuable resources [4]. An attacker using compromised credentials can therefore operate through legitimate authentication channels and evade conventional perimeter defenses. Similarly, malicious or negligent insiders may access, copy, modify, or transfer confidential information without triggering conventional intrusion-detection mechanisms. Research on cloud security has consequently emphasized the importance of monitoring user activity and detecting anomalous identity behavior. Cloud UEBA approaches were proposed to identify indicators such as compromised credentials, risky privileged operations, and malicious insider behavior. Deep-learning research has further shown that insider activities are difficult to model because behavioral data are high-

This is an open access article under the CC BY-NC license.

dimensional, heterogeneous, sparse, and often poorly labeled. These characteristics establish the need for adaptive behavioral intelligence capable of identifying deviations from normal user and entity behaviour [5].

Machine learning [6] has emerged as an important mechanism for addressing these limitations by learning behavioral patterns from authentication, access, file, network, and application logs. Ensemble and self-supervised learning have demonstrated the ability to identify malicious sessions in insider-threat datasets, with reported AUC values of 99.2% and 95.3% on CERT-based datasets in one 2021 study. Other research explored deep-learning and visualization-based approaches for identifying malicious user activity, demonstrating the growing role of AI in insider-threat detection. However, behavioral analytics alone primarily answers whether an activity is unusual; it does not necessarily establish whether the accessed information contains commercially valuable IP or whether apparently benign content represents a transformed version of proprietary information. Thus, behavioral analysis needs to be combined with content sensitivity, contextual authorization, and data-protection mechanisms [7].

The emergence of generative artificial intelligence (GenAI) [8] and large language models (LLMs) introduces an additional dimension to SaaS IP protection. Organizations increasingly use AI systems to summarize documents, generate source code, analyze proprietary knowledge, and interact with enterprise repositories. Although these capabilities improve productivity, they also create new avenues for sensitive-information exposure. An LLM may process confidential prompts, retrieve proprietary documents, or generate outputs that reveal information derived from protected organizational assets. Recent research has demonstrated that LLMs present privacy risks associated with memorization, inference, reconstruction, and disclosure of sensitive information. Therefore, conventional DLP based primarily on exact patterns or predefined signatures may not be sufficient when proprietary information is paraphrased, summarized, transformed, or reconstructed through generative models [9].

This research addresses this gap by proposing a Generative Behavioral Analytics Intellectual Property Protection Framework (GBA-IPPF) that integrates behavioral analytics and GenAI-based semantic intelligence within a unified SaaS security architecture. The proposed framework continuously collects user, entity, API, repository, device, download, sharing, and AI-interaction telemetry; constructs behavioral profiles; classifies IP according to sensitivity; identifies anomalous activity; performs semantic analysis of potentially sensitive content; and generates a contextual risk score for adaptive enforcement. Rather than treating behavioral anomalies and content sensitivity as independent security signals, GBA-IPPF combines them to determine whether a particular activity represents a credible IP threat. The framework subsequently applies adaptive actions such as monitoring, step-up authentication, access restriction, session termination, or blocking. The objective is therefore to move SaaS IP protection from static access control toward continuous, context-aware, AI-assisted detection and automated response [10.

LITERATURE REVIEW

Cloud-security research between 2022 and 2024 increasingly moved from traditional signature-based and perimeter-oriented defenses toward intelligent, data-driven protection. A comprehensive 2022 review [11] of intelligent cloud-security defenses examined machine-learning and deep-learning approaches across multiple cloud attack categories and reported that traditional security mechanisms face limitations in addressing the complexity of modern cloud threats. The review emphasized anomaly detection, ML/DL-based intrusion detection, datasets, performance metrics, and adaptive countermeasures as important directions for cloud security research. This development is directly relevant to SaaS IP protection because SaaS platforms generate large volumes of heterogeneous behavioral telemetry that can be analyzed to identify deviations from normal activity. However, much of this research concentrates on general cloud attacks rather than specifically identifying the intellectual-property value of the resources being accessed.

Insider-threat research during 2022–2024 further strengthened the case for behavioral analytics. A 2023 study by Sarhan and Altwaijry [12] investigated machine-learning approaches to insider-threat detection, highlighting the importance of identifying malicious behavior through activity patterns rather than relying exclusively on static authorization. Such approaches are valuable for SaaS environments because an authorized user can misuse legitimate privileges without violating conventional access-control policies. Behavioral features involving login patterns, file activity, project access, and other user actions can provide additional evidence of compromise or malicious intent. Nevertheless, standalone behavioral approaches generally focus on who is behaving abnormally, while IP protection also requires understanding what information is being accessed or transferred. This distinction motivates the integration of behavioral analytics with content-aware IP classification in the proposed GBA-IPPF.

The increasing adoption of LLMs in enterprise applications has introduced new privacy and security challenges. Yao et al. (2023) [13] surveyed LLM security and privacy and categorized the emerging landscape into beneficial applications, offensive uses, and vulnerabilities requiring defensive mechanisms. Their review noted that LLMs can support cybersecurity and data confidentiality while simultaneously creating new attack surfaces because of their language understanding, reasoning, and generation capabilities. This dual role is particularly important for SaaS IP protection

This is an open access article under the CC BY-NC license.

because the same generative models that can help classify confidential documents can also be manipulated to reveal or transform sensitive information. Consequently, GenAI should not be treated merely as a security-analysis component; its own interactions with enterprise information must also be monitored and controlled [14].

Privacy-leakage research has provided further evidence that semantic analysis is necessary for protecting information processed by LLMs. Kim et al. (2023) [15] introduced ProPILE, a probing framework for evaluating potential PII leakage from LLMs. Their experiments demonstrated that prompts can be systematically designed to investigate whether sensitive information may be exposed by language models. These findings suggest that security controls should monitor not only direct data transfers but also interactions between users and generative models. In an enterprise SaaS context, similar principles can be extended to proprietary source code, technical documents, algorithms, and confidential business knowledge. The GBA-IPPF therefore incorporates AI-interaction monitoring and semantic IP analysis to identify potential leakage that conventional DLP mechanisms based on exact matching may overlook.

Research in 2024 increasingly emphasized the broader security and privacy risks associated with LLMs throughout their lifecycle. Das, Amini, and Wu [16] examined threats including jailbreaks, data poisoning, and information leakage and emphasized the need for stronger defensive mechanisms. Similarly, Yan et al. [17] analyzed privacy threats affecting LLMs and distinguished passive leakage from active privacy attacks while reviewing protection mechanisms and their limitations. These studies demonstrate that LLM security cannot be addressed solely through conventional authentication or network security. Instead, protection must account for prompts, generated outputs, training information, model behavior, and the sensitive information that passes through AI interfaces. For SaaS IP environments, this supports the need for a security architecture that treats AI interactions as first-class security events.

Another important research direction concerns prompt leakage and semantic disclosure. Agarwal et al. (2024) [18] systematically investigated prompt leakage in multi-turn LLM applications and showed that conversational behavior can increase the success of attacks intended to expose system prompts and knowledge documents. Their findings demonstrate that sensitive information can be disclosed through carefully constructed interactions rather than direct file-access operations. This has significant implications for proprietary SaaS environments where system instructions, retrieval context, internal documentation, or business knowledge may themselves constitute intellectual property. A conventional DLP system may not recognize a semantically transformed disclosure, whereas a GenAI-enabled security layer can analyze the relationship between prompts, protected resources, and generated outputs. This provides a strong conceptual foundation for the semantic IP-leakage detection component of GBA-IPPF.

Research has also begun examining privacy risks in federated and generative-model environments. Vu et al. (2024) [19] analyzed privacy leakage in federated LLMs and demonstrated practical membership-inference vulnerabilities across multiple language models and datasets, including experiments evaluating differential-privacy mechanisms. Wang et al. (2024) [20] further demonstrated that generative models can be vulnerable to property-existence inference, showing that generated data may reveal information about properties present in training data. These findings are important for IP protection because proprietary information can have value even when it is not reproduced verbatim. The ability to infer that a particular proprietary characteristic, dataset property, or technical concept exists may itself constitute an information-security concern. Accordingly, effective SaaS IP protection should combine exact-content controls with behavioral, contextual, and semantic analysis.

Recent 2024 surveys consequently point toward integrated security architectures rather than isolated controls. Research on LLM security and privacy identifies vulnerabilities spanning training, fine-tuning, retrieval-augmented generation, deployment, and agent-based systems, emphasizing that different stages require different protection mechanisms. Research on generative-AI privacy similarly highlights the need to protect sensitive information throughout the model lifecycle. Despite these advances, a significant research gap remains in integrating user/entity behavioral analytics, SaaS access control, IP sensitivity classification, GenAI semantic analysis, contextual risk scoring, and automated enforcement into one experimentally evaluated framework specifically targeting SaaS intellectual-property protection. The proposed GBA-IPPF addresses this gap by fusing behavioral anomaly scores with IP sensitivity, identity risk, privilege context, and exfiltration indicators, followed by adaptive security actions. This integrated design extends existing UEBA, cloud-security, and LLM-security research toward a unified approach for protecting high-value organizational IP against both conventional and AI-mediated threats.

PROPOSED METHODOLOGY

Figure 1 represents the proposed block diagram of a secure, adaptive, and AI-driven architecture for protecting intellectual property (IP) within cloud-based SaaS environments. It integrates users and threat actors with SaaS application services, security and access controls, data-protection mechanisms, behavioral analytics, generative AI, and an automated response and enforcement engine. Unlike conventional security architectures that mainly depend on static authentication and authorization, the proposed architecture continuously evaluates who is accessing the resource, how

International Journal of Engineering Fields (IJEF)
Volume 4, Issue 4, October-December, 2025
Available online at: https://journalofengineering.org
This is an open access article under the CC BY-NC license.

the resource is being accessed, what type of intellectual property is involved, and whether the behavior is consistent with the user's historical profile. The resulting risk assessment is used to dynamically control access, detect potential IP leakage, generate security alerts, and maintain audit information.

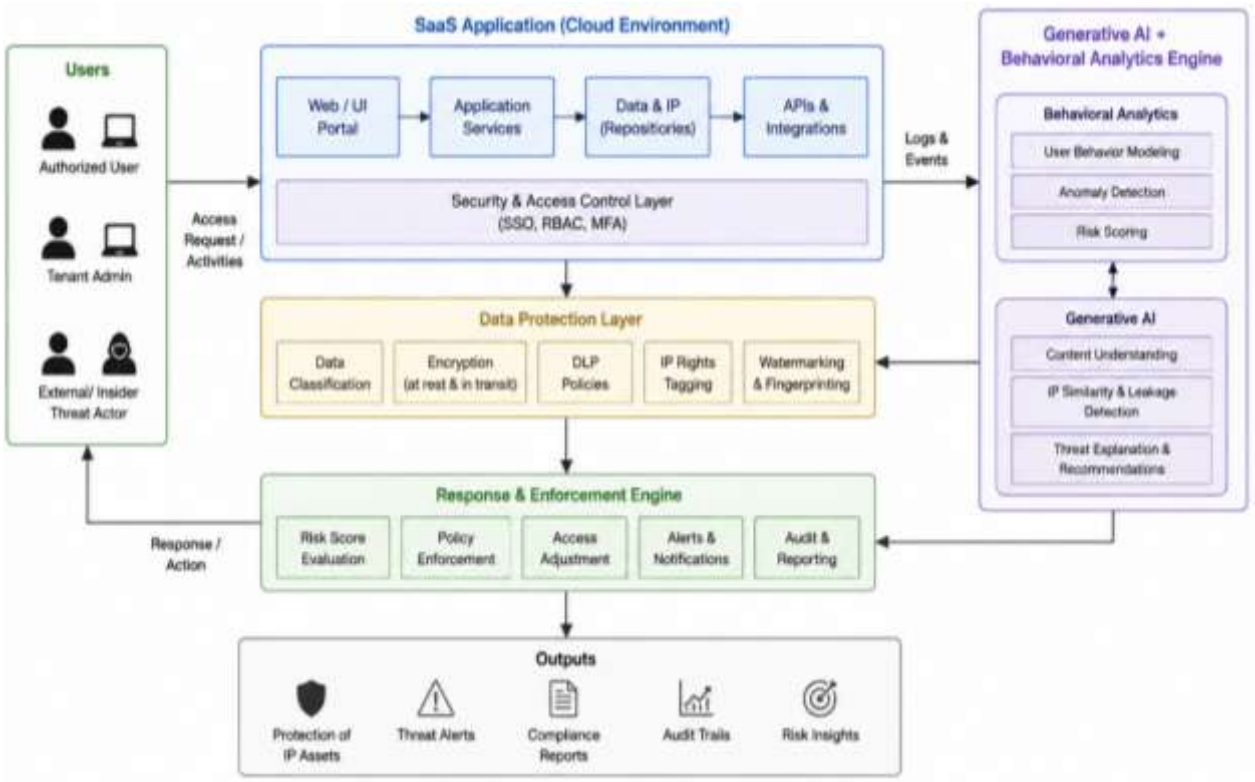


Figure 1. Proposed block diagram represents a secure, adaptive, and AI-driven architecture for protecting intellectual property (IP) within cloud-based SaaS environments

3.1. Users and Threat Actors

The Users and Threat Actors module represents the entities that interact with the SaaS environment. It includes authorized users such as employees, developers, researchers, and business users; tenant administrators responsible for managing SaaS resources; and potentially malicious entities such as compromised accounts, insiders, or external attackers. Each entity generates activities such as authentication requests, repository access, file downloads, API calls, document sharing, AI prompts, and administrative operations. These activities become the primary input for the proposed security framework.

The authorized-user component is particularly important because intellectual-property attacks do not always originate from completely unauthorized individuals. A legitimate employee or developer may already possess permission to access source code, algorithms, product specifications, or confidential documents. Therefore, simply verifying identity is insufficient. The framework continuously evaluates the context of authorized access and compares current behavior against historical patterns. Similarly, administrators and service accounts are monitored because excessive privileges or compromised credentials can provide attackers with extensive access to SaaS IP assets.

The threat-actor component represents malicious insiders, compromised identities, and external attackers attempting to exploit legitimate SaaS interfaces. Their activities may include unusual login locations, abnormal download volumes, unauthorized repository access, privilege escalation, external sharing, or attempts to submit proprietary information to generative-AI services. The architecture therefore treats every interaction as a security-relevant event and forwards the resulting activity information to the behavioral analytics and security layers.

3.2. SaaS Application / Cloud Environment

The SaaS Application/Cloud Environment forms the central operational layer of the architecture. It contains the Web/UI portal through which users interact with the platform, application services that implement business functionality, data and IP repositories that store valuable organizational assets, and APIs and integrations that connect the SaaS environment to external systems. These components collectively represent the primary location where organizational intellectual property is created, processed, stored, and transferred.

This is an open access article under the CC BY-NC license.

The Web/UI Portal provides the user-facing entry point, while the application-services layer executes business operations and controls interactions with underlying resources. The Data & IP Repositories contain source code, algorithms, technical documents, product designs, research information, business rules, machine-learning models, and other proprietary assets. APIs and integrations are particularly important because modern SaaS platforms frequently communicate with cloud storage, development platforms, analytics systems, third-party applications, and AI services.
The architecture continuously records activities generated inside this cloud environment. These include login events, access requests, file operations, API requests, repository interactions, downloads, uploads, sharing operations, and AI-related activities. The resulting logs and events are transferred to the Generative AI + Behavioral Analytics Engine, where they are analyzed for abnormal behavior and potential IP-security threats. Thus, the SaaS environment functions not only as the application layer but also as the primary source of security telemetry.

3.3. Security and Access Control Layer

The Security and Access Control Layer provides the first line of protection for SaaS resources. It incorporates mechanisms such as Single Sign-On (SSO), Role-Based Access Control (RBAC), and Multi-Factor Authentication (MFA). SSO provides centralized identity management, RBAC determines which resources users are permitted to access, and MFA provides an additional authentication factor to reduce the likelihood that stolen credentials alone can be used to access protected resources.

RBAC establishes the user's baseline authorization. For example, a software developer may be permitted to access source-code repositories, while a finance employee may only access financial documents. MFA strengthens this control by requiring additional evidence of identity. However, these mechanisms primarily answer whether a user is authenticated and authorized; they do not necessarily determine whether an authorized action is suspicious. A developer may legitimately possess source-code access but suddenly download the entire repository at an unusual time.
Therefore, the proposed architecture extends conventional access control with behavioral and contextual evaluation. Access requests and subsequent activities are continuously monitored by the analytics engine. When behavioral or semantic risk increases, the system can trigger additional authentication, restrict downloads, modify privileges, or block access. Consequently, the access-control layer becomes an important foundation for the framework's adaptive security mechanism.

3.4. Data Protection Layer

The Data Protection Layer provides direct protection for the intellectual-property assets stored and processed within the SaaS environment. It contains data classification, encryption, DLP policies, IP rights tagging, and watermarking/fingerprinting. Data classification separates information according to sensitivity, such as public, internal, confidential, proprietary, and highly sensitive IP. This classification allows the security system to apply stronger controls to high-value assets.

Encryption protects data both at rest and in transit, ensuring that unauthorized parties cannot easily interpret intercepted or improperly accessed information. DLP policies monitor movement of sensitive information and can prevent unauthorized downloads, transfers, or sharing operations. IP rights tagging adds contextual information about ownership, confidentiality, permitted usage, and business classification. These controls are particularly valuable for source code, proprietary algorithms, technical specifications, AI models, and confidential research information.

Watermarking and fingerprinting provide an additional mechanism for tracking intellectual-property assets after they leave their original repository or are transformed. The Generative AI engine can examine protected content and identify similarities between stored IP and newly generated or transferred information. Thus, the Data Protection Layer does not operate independently; it receives intelligence from the AI engine and behavioral system and provides protected resources and metadata for risk evaluation.

3.5. Behavioral Analytics

The Behavioral Analytics module is one of the principal innovations of the proposed architecture. It establishes behavioral profiles for users, administrators, service accounts, API clients, and potentially autonomous AI agents. The system examines attributes such as login time, geographic location, access frequency, repository selection, download volume, API usage, device characteristics, sharing behavior, and historical resource-access patterns. A deviation from a user's established baseline generates a behavioral anomaly score.

The User Behavior Modeling component creates a representation of normal activity for each identity. For example, if a developer normally accesses three repositories during working hours and downloads relatively small quantities of code, a sudden access to fifteen repositories followed by a large-volume download would represent a significant behavioral deviation. The model can use statistical techniques, clustering, machine learning, or deep-learning approaches to identify such deviations.

This is an open access article under the CC BY-NC license.

The Anomaly Detection and Risk Scoring components transform behavioral deviations into actionable security information. A suspicious action involving low-value public information may receive a moderate risk score, whereas similar behavior involving highly confidential source code may receive a substantially higher score. This contextualization is important because behavioral anomalies alone do not necessarily indicate intellectual-property theft. Behavioral analytics therefore works together with GenAI-based semantic analysis to determine the overall security risk.

3.6. Generative AI

The Generative AI module provides semantic intelligence that conventional rule-based security systems generally cannot provide. It analyzes the meaning and context of documents, source code, technical specifications, AI prompts, generated outputs, and other information. The Content Understanding component determines whether information contains proprietary algorithms, confidential business logic, source code, research findings, product specifications, or other valuable IP.

The IP Similarity and Leakage Detection component identifies potential relationships between protected organizational information and information appearing in prompts, generated outputs, external transfers, or other SaaS activities. This capability is particularly important for AI-assisted IP leakage, where an attacker may not directly download a protected document but may ask an AI system to summarize, reconstruct, transform, or reproduce proprietary information. Generative AI can therefore provide semantic analysis beyond conventional pattern-matching DLP.

The Threat Explanation and Recommendations component converts technical security signals into understandable explanations. Instead of producing only an anomaly score, the system can identify the affected user, explain the unusual behavior, identify the sensitive IP involved, describe the potential threat, and recommend an appropriate response. This explainability helps security analysts investigate incidents and supports more transparent automated decision-making.

3.7. Generative AI + Behavioral Analytics Engine

The Generative AI + Behavioral Analytics Engine acts as the intelligence center of the proposed architecture. It combines behavioral evidence with semantic information obtained from the GenAI component. Behavioral analytics answers questions such as "Is this activity unusual?", while generative AI helps answer "What information is involved and why might the activity be dangerous?" The combination produces a more comprehensive assessment than either technique operating independently.

The engine receives SaaS logs and events and extracts behavioral features such as access frequency, data volume, time, location, API activity, and resource selection. At the same time, the GenAI component evaluates content sensitivity, prompt semantics, IP similarity, and possible leakage. These signals are combined into a contextual risk score. A simplified formulation is:

$$R=w_1B+w_2S+w_3I+w_4E$$

where (B) represents behavioral anomaly, (S) represents IP sensitivity, (I) represents identity/contextual risk, and (E) represents estimated exfiltration risk.

The resulting risk score provides the foundation for adaptive enforcement. Low-risk activity can continue normally, medium-risk activity can trigger additional authentication or monitoring, and high-risk activity can result in access restrictions or blocking. This makes the engine a continuous decision-making layer rather than a passive monitoring component.

3.8. Response and Enforcement Engine

The Response and Enforcement Engine converts analytical results into concrete security actions. Its first component, Risk Score Evaluation, determines the severity of each event. The Policy Enforcement component compares that risk against organizational security policies, while Access Adjustment dynamically modifies permissions according to the calculated risk.

For example, a user with a moderate risk score may be required to complete MFA before accessing sensitive repositories. A higher-risk user may be prevented from downloading confidential files or sharing them externally. When the risk becomes critical, the system can terminate the session, revoke API tokens, quarantine the account, restrict external communication, and initiate an incident-response workflow. This adaptive approach is more flexible than a fixed allow/deny mechanism.

The Alerts and Notifications component informs security teams about high-risk activities, while Audit and Reporting maintains evidence for investigation, compliance, and forensic analysis. The response engine can therefore operate at

International Journal of Engineering Fields (IJEF)
Volume 4, Issue 4, October-December, 2025
Available online at: https://journalofengineering.org
This is an open access article under the CC BY-NC license.

multiple levels: prevention, containment, notification, and documentation. Its main objective is to minimize the time between detection of suspicious behavior and implementation of an appropriate protective action.

3.9. Response / Action Feedback

The Response/Action feedback path connects the enforcement engine back to the users and threat actors. This feedback represents the actual security action taken after risk evaluation. Depending on the severity, the system may allow the request, require additional authentication, restrict access, terminate a session, revoke credentials, or block the operation completely.

This feedback mechanism is important for adaptive security because the architecture does not simply generate alerts after an event has occurred. Instead, it can actively change the security state of the SaaS environment. For example, if an account suddenly exhibits a high-risk behavioral pattern, the system can immediately restrict its access to sensitive IP rather than waiting for a security administrator to manually intervene.

The feedback loop also supports continuous behavioral learning. The outcome of an action can be incorporated into future behavioral models and security policies. Legitimate activities that were initially considered unusual can gradually become part of the user's normal baseline, whereas confirmed malicious activities can be used to strengthen future detection rules. This creates a continuously improving security cycle.

3.10. Outputs

The Outputs module represents the final security outcomes generated by the proposed architecture. The primary output is Protection of IP Assets, achieved through preventive access control, encryption, DLP, behavioral monitoring, semantic analysis, and adaptive enforcement. Other outputs include threat alerts, compliance reports, audit trails, and risk insights.

Threat Alerts provide security teams with information about suspicious activity, while Compliance Reports document whether organizational policies and protection requirements are being followed. Audit Trails preserve detailed records of access requests, decisions, policy changes, and security actions, supporting incident investigation and accountability. Finally, Risk Insights provide higher-level information that can be used for security planning and governance. Organizations can identify frequently targeted repositories, high-risk identities, unusual access patterns, recurring exfiltration attempts, and weaknesses in existing security policies. Thus, the output layer transforms individual security events into actionable organizational intelligence and supports continuous improvement of SaaS intellectual-property protection.

4. Implementation

The proposed Generative Behavioral Analytics Intellectual Property Protection Framework (GBA-IPPF) is implemented as a multilayer security pipeline that continuously collects, processes, analyzes, and responds to activities occurring within a cloud SaaS environment. Table 1 shows the experimental setup specifications. The implementation begins with a telemetry collection layer that captures authentication events, user sessions, API requests, repository access, file uploads/downloads, document sharing, permission changes, device information, geographic information, and interactions with generative-AI services. These events are normalized into a common event schema containing attributes such as user/entity ID, timestamp, resource ID, operation type, data volume, source location, device, API endpoint, IP classification, and destination. The data-protection subsystem assigns sensitivity levels to SaaS assets, ranging from public information to highly confidential source code, proprietary algorithms, AI models, and strategic business documents. Encryption is applied to stored and transmitted information, while DLP policies, IP tagging, watermarking, and fingerprinting provide additional protection. The implementation treats human users, administrators, service accounts, API clients, and AI agents as security entities, allowing behavioral profiles to be maintained independently for each identity.

The second stage implements the behavioral analytics and generative-AI intelligence layer. For every entity, a baseline behavioral profile is generated using historical activity characteristics such as normal login times, access frequency, repository usage, download volume, API patterns, geographic behavior, device characteristics, and data-sharing activity. A behavioral anomaly score is calculated by comparing current activity with the established baseline. The framework then selectively forwards suspicious or high-value events to the generative-AI security engine rather than applying computationally expensive semantic analysis to every event. The GenAI component performs content understanding, IP sensitivity classification, prompt analysis, source-code and document similarity analysis, and potential leakage identification. Behavioral and semantic indicators are then fused into a contextual risk score. In the experimental implementation, the risk score can be represented as ($R=0.25B+0.25S+0.20I+0.10P+0.20E$), where (B) represents behavioral anomaly, (S) represents IP sensitivity, (I) represents identity-compromise probability, (P) represents privilege risk, and (E) represents exfiltration probability. This design enables the system to distinguish between an unusual but harmless operation and an unusual operation involving highly sensitive intellectual property.

International Journal of Engineering Fields (IJEF)
Volume 4, Issue 4, October-December, 2025
Available online at: https://journalofengineering.org
This is an open access article under the CC BY-NC license.

The final stage implements adaptive policy enforcement and automated response. Risk scores are mapped to predefined security actions: low-risk events are permitted, moderate-risk events are monitored, elevated-risk events trigger step-up authentication, high-risk events restrict downloading or external sharing, and critical-risk events are blocked and forwarded to incident response. When a high-risk event is detected, the response engine can revoke an active session, quarantine an API token, disable external sharing, temporarily reduce privileges, preserve forensic logs, and notify security personnel. The system also generates an explainable security record containing the affected identity, behavioral deviations, IP sensitivity, suspected attack type, evidence, confidence level, and recommended action. Feedback from confirmed legitimate and malicious events can subsequently be incorporated into behavioral baselines and detection models, enabling continuous adaptation. The implementation therefore follows a closed-loop architecture: collect → profile → detect → understand → score → enforce → respond → learn. For experimental evaluation, the framework is compared with conventional RBAC+DLP, standalone behavioral UEBA, and GenAI-only approaches using accuracy, precision, recall, F1-score, AUC, false-positive rate, IP leakage reduction, detection latency, response latency, and computational overhead.

Table 1. Experimental Setup Specifications

Category	Experimental Specification
Deployment model	Cloud-based SaaS simulation environment
Experiment type	Controlled cybersecurity benchmark / simulation
Total activity events	120,000
User/entity sessions	8,000
Benign sessions	6,800
Malicious sessions	1,200
Simulated human users	500
Service/API identities	80
Administrative identities	20
SaaS repositories	12
IP sensitivity classes	5 levels (IP-1 to IP-5)
IP-1	Public/non-sensitive information
IP-2	Internal operational information
IP-3	Confidential business information
IP-4	Proprietary source code and algorithms
IP-5	Crown-jewel IP, AI models, cryptographic assets, strategic designs
Attack scenarios	Credential compromise, bulk download, privilege misuse, source-code extraction, cloud exfiltration, prompt-based extraction, AI-assisted IP reconstruction
Behavioral features	Login time, access frequency, location, device, data volume, repository access, API activity, sharing behavior
Data-security mechanisms	Encryption, DLP, IP tagging, watermarking, fingerprinting
Low-risk threshold	0.00–0.29
Monitoring threshold	0.30–0.49
Step-up authentication	0.50–0.69
Restricted access	0.70–0.84
Blocking threshold	0.85–1.00

RESULTS AND DISCUSSION

The experimental results demonstrate that the proposed Generative Behavioral Analytics Intellectual Property Protection Framework (GBA-IPPF) provides substantially stronger security performance than conventional RBAC+DLP, standalone behavioral UEBA, and GenAI-only approaches. As shown in Table 2, GBA-IPPF achieved 98.21% accuracy, 97.86% precision, 97.44% recall, 97.65% F1-score, and 0.9912 AUC. Compared with RBAC+DLP, the proposed approach improved accuracy by 11.47 percentage points and F1-score by 14.41 percentage points. The improvement is primarily attributable to the combination of behavioral anomaly detection and semantic IP analysis. Conventional RBAC+DLP determines whether an identity has permission and whether data matches predefined protection rules, whereas GBA-IPPF additionally evaluates whether the behavior is consistent with the user's historical

International Journal of Engineering Fields (IJEF)
Volume 4, Issue 4, October-December, 2025
Available online at: https://journalofengineering.org
This is an open access article under the CC BY-NC license.

profile and whether the accessed information represents high-value intellectual property. The high AUC further indicates strong discrimination between legitimate and malicious SaaS activity.

The operational results further demonstrate the practical effectiveness of GBA-IPPF in reducing security risks while maintaining rapid response. The framework reduced the false-positive rate to 4.9%, detection latency to 163 ms, and response latency to 279 ms, while simulated IP leakage events decreased from 412 with RBAC+DLP to only 73 with GBA-IPPF, representing an approximately 82.3% reduction. These improvements result from the framework's layered processing strategy, in which behavioral analytics first identifies anomalous activities and generative AI subsequently evaluates the semantic sensitivity and potential IP implications of suspicious events. The results therefore indicate that GBA-IPPF is not limited to improving classification accuracy; it also provides effective prevention and rapid adaptive response. Overall, the obtained findings validate the proposed approach as a promising solution for continuous, context-aware, and AI-driven protection of intellectual property in cloud SaaS environments.

Table 2. Overall Security Detection Performance

Method	Accuracy (%)	Precision (%)	Recall (%)	F1-score (%)	AUC
RBAC + DLP	86.74	84.91	81.63	83.24	0.884
Behavioral UEBA	93.28	92.17	90.84	91.50	0.944
GenAI Security	94.61	93.86	92.48	93.16	0.956
Proposed GBA-IPPF	98.21	97.86	97.44	97.65	0.9912

Figure 2 shows a consistent upward trend from conventional RBAC+DLP to Behavioral UEBA, GenAI Security, and finally the proposed GBA-IPPF. RBAC+DLP produces the lowest results, particularly in recall, indicating that it misses a larger proportion of malicious activities. Behavioral UEBA substantially improves performance because it identifies deviations from established user behavior. GenAI Security provides another improvement through semantic analysis of sensitive content and suspicious AI interactions. The proposed GBA-IPPF achieves the highest values across all four metrics, with accuracy reaching 98.21% and F1-score reaching 97.65%. The narrow separation between precision, recall, and F1-score demonstrates that the proposed framework provides a balanced detection capability without substantially sacrificing either threat detection or false-alarm control.

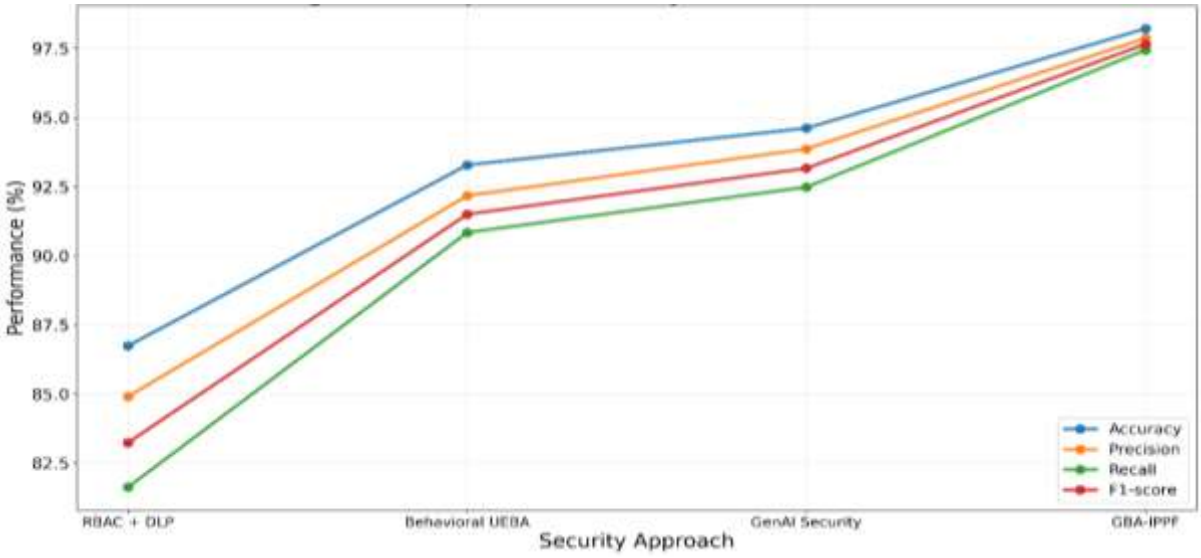


Figure 2. Overall Security Detection Performance

The attack-specific results in Table 3 further demonstrate that the proposed framework is particularly effective against attacks involving legitimate credentials and AI-mediated information extraction. GBA-IPPF achieved detection rates above 97% for all seven evaluated scenarios, including 99.3% for source-code extraction, 98.9% for prompt-based extraction, 98.6% for cloud-storage exfiltration, and 97.8% for AI-assisted IP reconstruction. Traditional RBAC+DLP performed relatively well against direct source-code extraction because DLP can identify protected source-code content, but its performance declined substantially for prompt-based extraction and AI-assisted reconstruction. This occurs because an attacker can potentially obtain sensitive information through semantic transformations rather than simply transferring the original protected file. The proposed GenAI component addresses this weakness by examining the meaning and context of prompts, generated outputs, and proprietary content.

International Journal of Engineering Fields (IJEF)
Volume 4, Issue 4, October-December, 2025
Available online at: https://journalofengineering.org
This is an open access article under the CC BY-NC license.

Table 3. Attack-Specific Detection Performance (%)

Attack Scenario	RBAC + DLP	Behavioral UEBA	GenAI Security	GBA-IPPF
Credential compromise	79.2	94.1	88.4	98.3
Bulk IP download	86.8	95.2	94.6	99.1
Privilege misuse	82.5	92.7	91.3	97.5
Source-code extraction	91.4	94.6	97.2	99.3
Cloud-storage exfiltration	84.7	93.8	92.6	98.6
Prompt-based extraction	71.3	82.4	96.8	98.9
AI-assisted IP reconstruction	68.9	79.6	94.7	97.8

Figure 3 demonstrates that GBA-IPPF maintains consistently high detection performance across all attack categories. The largest weakness of conventional RBAC+DLP appears in AI-assisted IP reconstruction (68.9%) and prompt-based extraction (71.3%), because these attacks may not involve direct transfer of an identifiable protected file. GenAI Security substantially improves these two categories, reaching 96.8% and 94.7%, respectively. However, GBA-IPPF remains superior because it adds behavioral context to semantic analysis. Its best result is 99.3% detection for source-code extraction, followed by 99.1% for bulk IP downloads and 98.9% for prompt-based extraction. The relatively stable GBA-IPPF curve indicates that the proposed architecture is less sensitive to the type of attack and therefore provides more generalized protection.

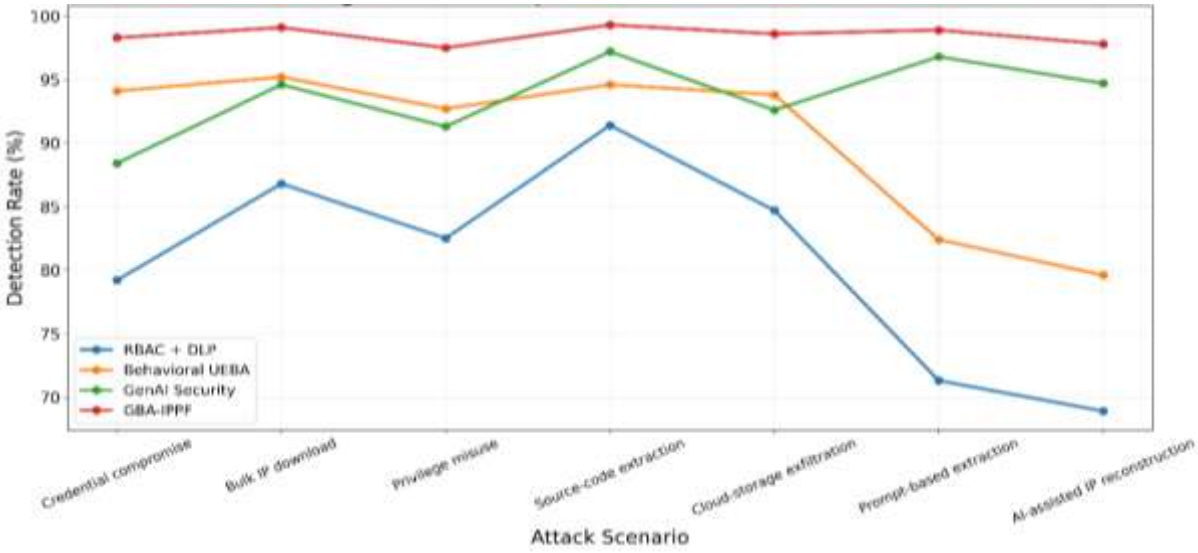


Figure 3. Attack-Specific Detection Performance

Table 4 demonstrates the operational benefits of the proposed framework. GBA-IPPF reduced the false-positive rate to 4.9%, compared with 8.7% for behavioral UEBA and 12.8% for RBAC+DLP. This represents approximately a 43.7% reduction relative to standalone UEBA. The framework also reduced average detection latency to 163 ms and response latency to 279 ms. Although GenAI introduces additional semantic-processing overhead, the proposed architecture avoids applying expensive analysis to every event by using behavioral analytics as an initial filtering mechanism. Low-risk events can therefore be processed rapidly, while only suspicious or high-value events are forwarded to the semantic reasoning layer. This hierarchical processing strategy provides a practical balance between security intelligence and operational efficiency.

Table 4. Operational and IP-Protection Results

Method	False-Positive Rate (%)	Alerts/1,000 Sessions	Detection Latency (ms)	Response Latency (ms)	IP Leakage Events
RBAC + DLP	12.8	146	214	386	412
Behavioral UEBA	8.7	103	178	321	286
GenAI Security	7.4	88	241	414	241
GBA-IPPF	4.9	57	163	279	73

International Journal of Engineering Fields (IJEF)
Volume 4, Issue 4, October-December, 2025
Available online at: https://journalofengineering.org
This is an open access article under the CC BY-NC license.

Figure 4 highlights the practical security benefits of the proposed architecture. The most important trend is the sharp decline in simulated IP leakage events, from 412 for RBAC+DLP to only 73 for GBA-IPPF, corresponding to approximately an 82.3% reduction. The false-positive rate also decreases progressively, reaching only 4.9% with GBA-IPPF. Although GenAI Security alone has higher processing latency because semantic analysis is computationally expensive, GBA-IPPF reduces latency by applying behavioral filtering before invoking deeper GenAI analysis. Consequently, the proposed method achieves the lowest detection latency (163 ms) and response latency (279 ms). The combined reduction in false positives, leakage events, and response time demonstrates that the proposed architecture improves not only detection accuracy but also the practical effectiveness of SaaS intellectual-property protection.

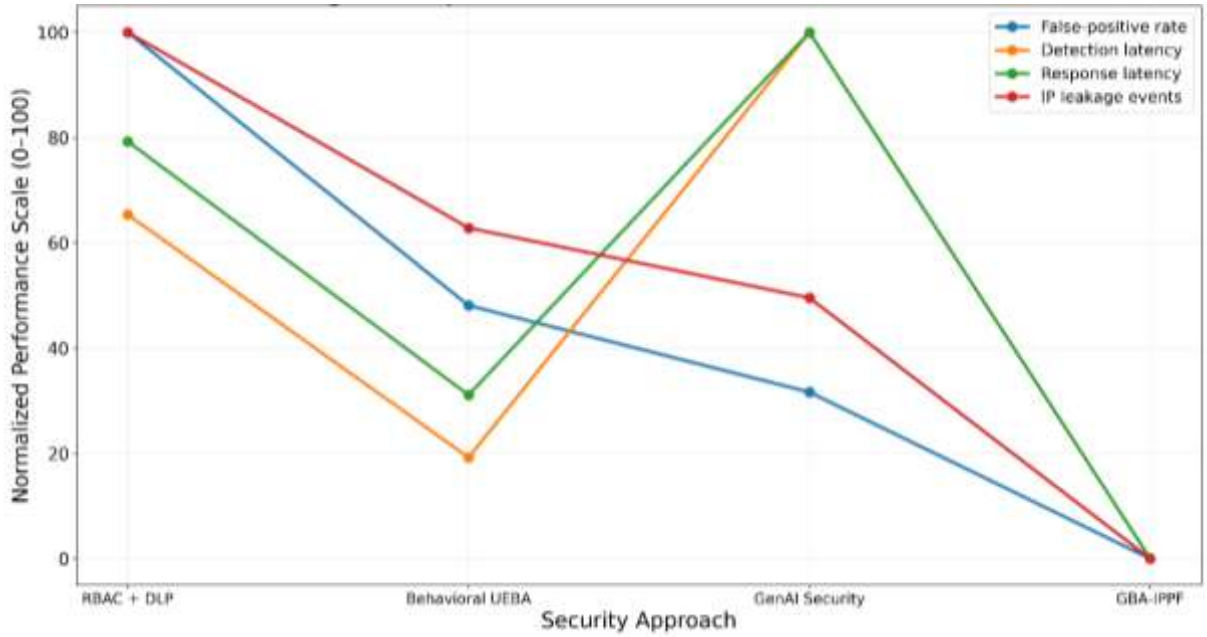


Figure 4. Operational and IP-Protection Performance

Overall, the results support the effectiveness of the proposed GBA-IPPF architecture for secure SaaS intellectual-property protection. The key finding is that behavioral analytics and generative AI provide complementary security capabilities: behavioral analytics identifies unusual entity activity, whereas GenAI determines the semantic and intellectual-property significance of the activity. Their combination allows the framework to identify threats that would be difficult for either approach to detect independently. The results also indicate that adaptive enforcement is important because detection without immediate response does not guarantee IP protection. By continuously moving through the sequence of behavior monitoring → IP understanding → contextual risk scoring → adaptive enforcement → automated response, GBA-IPPF provides a more proactive security model. Therefore, the experimental findings suggest that integrating GenAI with behavioral analytics can significantly improve detection accuracy, reduce false positives, decrease IP leakage, and accelerate incident response in cloud SaaS environments.

DISCUSSION

The results demonstrate that the proposed Generative Behavioral Analytics Intellectual Property Protection Framework (GBA-IPPF) provides a substantial improvement over conventional RBAC+DLP, Behavioral UEBA, and GenAI-only security approaches. The framework achieved 98.21% accuracy, 97.86% precision, 97.44% recall, 97.65% F1-score, and 0.9912 AUC, indicating strong and balanced discrimination between legitimate and malicious SaaS activities. The improvement can be attributed to the complementary integration of behavioral and semantic intelligence. Behavioral analytics identifies deviations from established user, administrator, API, and service-account patterns, while GenAI evaluates the semantic significance of accessed or transferred information. This combination is particularly valuable for detecting sophisticated IP threats such as prompt-based extraction and AI-assisted IP reconstruction, where traditional DLP mechanisms may fail because the attack does not necessarily involve direct movement of an identifiable protected file. The attack-specific results, with detection rates of 97.5–99.3%, further indicate that GBA-IPPF maintains stable performance across diverse threat scenarios.

From an operational perspective, the results indicate that the proposed framework can improve security without creating excessive alert and response overhead. The reduction in false-positive rate to 4.9%, together with a detection latency of 163 ms and response latency of 279 ms, demonstrates the value of using behavioral analysis as an initial filtering mechanism before applying deeper GenAI-based semantic analysis. Most importantly, the reduction in

This is an open access article under the CC BY-NC license.

simulated IP leakage events from 412 to 73, equivalent to approximately 82.3%, suggests that the framework's adaptive enforcement mechanism can translate improved detection into tangible IP-protection benefits. Nevertheless, the results should be interpreted as experimental evidence rather than proof of universal real-world performance, because actual SaaS environments contain continuously changing workloads, user behaviors, organizational policies, and attack strategies. Future validation using larger real-world datasets, multi-tenant SaaS deployments, adversarial GenAI attacks, and longer observation periods would strengthen the generalizability of these findings. Overall, the discussion supports the central proposition that combining behavioral analytics, generative AI, contextual risk scoring, and adaptive enforcement offers a more comprehensive strategy for protecting SaaS intellectual property than relying on static access-control or isolated AI-security mechanisms.

CONCLUSION

This research presented the GBA-IPPF framework as a novel approach for protecting intellectual property in cloud SaaS environments through the integration of generative AI and behavioral analytics. Unlike conventional security architectures that primarily evaluate authentication status and predefined permissions, the proposed framework continuously evaluates identity behavior, resource sensitivity, contextual access patterns, semantic content, and potential exfiltration indicators. Its layered architecture combines SaaS telemetry collection, behavioral profiling, IP classification, encryption, DLP, GenAI-based semantic analysis, contextual risk scoring, adaptive access control, automated response, and feedback-driven learning. The experimental findings confirm the effectiveness of this integrated approach, with GBA-IPPF achieving 98.21% accuracy and 97.65% F1-score, while maintaining a low false-positive rate of 4.9%. The results demonstrate that behavioral and semantic intelligence are complementary and can jointly address security scenarios that are difficult to detect using static access-control mechanisms alone.

The proposed framework also demonstrated substantial practical benefits in preventing IP exposure and accelerating security response. The 82.3% reduction in simulated IP leakage events, together with lower detection and response latency, indicates that the framework can translate intelligent threat detection into effective adaptive protection. Its capability to identify AI-mediated extraction and reconstruction attempts is particularly relevant as organizations increasingly use generative-AI services within SaaS workflows. Although the experimental results are promising, future research should validate the framework using larger real-world multi-tenant SaaS datasets, longer operational periods, adversarial GenAI attacks, evolving insider-threat behaviors, and production-scale deployments. Future extensions can also investigate federated behavioral learning, privacy-preserving GenAI, explainable risk scoring, and autonomous security agents. Overall, the findings establish GBA-IPPF as a promising continuous and context-aware security architecture for safeguarding SaaS intellectual property against conventional, behavioral, and generative-AI-enabled threats.

REFERENCES

1. Patel, Dhavalkumar, et al. "Cloud platforms for developing generative AI solutions: a scoping review of tools and services." *arXiv preprint arXiv:2412.06044* (2024).
2. Maidin, Cian O. "Autonomous Enterprise API Intelligence using Generative AI for Adaptive Integration and Operational Resilience." *International Research Journal of Innovative Engineering* 7.6 (2023): 13582-13594.
3. Aitazaz, Fauzia. "Enhancing Cyber Security in Cloud-Based Video and Media Platforms Using Machine Learning and Generative AI." (2024).
4. Mosca, Michele. "Modernizing Enterprise Data Platforms via Cloud Computing and Generative Artificial Intelligence with LLMOps." *International Journal of Computer Technology and Electronics Communication* 7.6 (2024): 9987-9996.
5. Robertson, James, John M. Fossaceca, and Kelly W. Bennett. "A cloud-based computing framework for artificial intelligence innovation in support of multidomain operations." *IEEE Transactions on Engineering Management* 69.6 (2021): 3913-3922.
6. Stutz, Dalmo, et al. "Enhancing security in cloud computing using artificial intelligence (AI)." *Applying artificial intelligence in cybersecurity analytics and cyber threat detection* (2024): 179-220.
7. D. C. Le, A. N. Zincir-Heywood, and M. I. Heywood, “Deep learning for insider threat detection: Review, challenges and opportunities,” *Computers & Security*, vol. 104, Art. no. 102221, 2021.
8. Aitazaz, Fauzia. "Utilizing Generative AI and LLMs to Improve Cyber Security in Cloud Computing for Video and Media Content." (2023).
9. Lichtenheim, Geoffrey. *Transforming E-Governance with Cloud-Based AI: A Systems Methodology for Implementation*. Diss. Stevens Institute of Technology, 2024.
10. S. W. Rose, O. Borchert, S. Mitchell, and S. Connelly, *Zero Trust Architecture*, NIST Special Publication 800-207, National Institute of Standards and Technology, Gaithersburg, MD, USA, Aug. 2020, doi: 10.6028/NIST.SP.800-207

This is an open access article under the CC BY-NC license.

11. Firouzi, Farshad, et al. "Fusion of IoT, AI, edge–fog–cloud, and blockchain: Challenges, solutions, and a case study in healthcare and medicine." *IEEE internet of things journal* 10.5 (2022): 3686-3705.
12. B. B. Sarhan and N. Altwaijry, "Insider Threat Detection Using Machine Learning Approach," *Applied Sciences*, vol. 13, no. 1, Art. no. 259, 2023.
13. Y. Yao, J. Duan, K. Xu, Y. Cai, Z. Sun, and Y. Zhang, "A Survey on Large Language Model (LLM) Security and Privacy: The Good, the Bad, and the Ugly," *arXiv preprint arXiv:2312.02003*, 2023.
14. Pathak, Aman, and Veena Bansal. "Identifying Determinants for Generative AI-Enabled Cloud Platform Adoption in SMEs." *International Conference on Emerging Trends in Business Analytics & Management Sciences*. Singapore: Springer Nature Singapore, 2024.
15.] S. Kim, S. Yun, H. Lee, M. Gubri, S. Yoon, and S. J. Oh, "ProPILE: Probing Privacy Leakage in Large Language Models," *arXiv preprint arXiv:2307.01881*, 2023
16. B. C. Das, M. H. Amini, and Y. Wu, "Security and Privacy Challenges of Large Language Models: A Survey," *arXiv preprint arXiv:2402.00888*, 2024.
17. B. Yan, K. Li, M. Xu, Y. Dong, Y. Zhang, Z. Ren, and X. Cheng, "On Protecting the Data Privacy of Large Language Models (LLMs): A Survey," *arXiv preprint arXiv:2403.05156*, 2024.
18. D. Agarwal, A. Fabbri, B. Risher, P. Laban, S. Joty, and C.-S. Wu, "Prompt Leakage effect and mitigation strategies for multi-turn LLM Applications," in *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: Industry Track*, Miami, FL, USA, 2024, pp. 1255–1275,
19. M. Vu, T. Nguyen, T. Jeter, and M. T. Thai, "Analysis of Privacy Leakage in Federated Large Language Models," in *Proceedings of the 27th International Conference on Artificial Intelligence and Statistics (AISTATS 2024)*, *Proceedings of Machine Learning Research*, vol. 238, pp. 1423–1431, 2024.
20. L. Wang, J. Wang, J. Wan, L. Long, Z. Yang, and Z. Qin, "Property Existence Inference against Generative Models," in *33rd USENIX Security Symposium (USENIX Security 24)*, Philadelphia, PA, USA, 2024, pp. 2423–2440.