

Unified Framework for Securing Cloud-Native Storage: Approach for Detecting and Mitigating Multi-Cloud Bucket Misconfigurations

Sanat Talwar

Dept. of Security

Electronic Arts, Inc.

Austin, Texas sanattalwar1994@gmail.com

Abstract—

Misconfigurations in cloud-native storage buckets across multi-cloud environments pose substantial security risks. These vulnerabilities can result in unauthorized access, data breaches, regulatory violations, and considerable financial and reputational consequences for businesses. The complexity of securing cloud storage is heightened by the variety of security models, access control frameworks, and API architectures among prominent cloud service providers such as AWS, Google Cloud, and Azure. This diversity complicates organizations' efforts to implement consistent and effective security protocols, leaving cloud storage resources at risk of misconfigurations that can be challenging to detect and rectify[4].

This paper presents a comprehensive, automated framework engineered to identify, evaluate, and remediate misconfigurations in cloud-native storage services within multi-cloud environments. The proposed framework utilizes state-of-the-art cloud-native tools, automated scanning techniques, and real-time risk assessment functionalities to efficiently detect vulnerable storage buckets, ascertain their risk levels, and execute timely remediation strategies. By integrating external threat intelligence sources, including both public and proprietary feeds, the framework enhances the identification of potential threats, including anomalous activities or known vulnerabilities associated with misconfigured storage resources[2].

Beyond threat intelligence integration, the framework employs sophisticated anomaly detection algorithms that scrutinize cloud storage configurations and access patterns to pinpoint deviations from standard operational behavior. These algorithms are essential for recognizing subtle misconfigurations that may otherwise remain undetected. Additionally, the framework encompasses policy enforcement tools that empower organizations to automatically define and uphold cloud security policies, ensuring that all storage resources adhere to established security guidelines and standards[6].

Experimental evaluations across diverse multi-cloud environments demonstrate significant enhancements in detection accuracy, risk assessment precision, and scalability. The framework effectively alleviates the manual burden typically associated with conventional cloud security management processes, allowing security teams to concentrate on high-priority tasks instead of dedicating time to routine checks and remediation activities. The results underscore the framework's ability to automate misconfiguration identification, prioritize critical risks based on potential impact, and maintain ongoing security compliance in real time, thereby addressing a significant gap in the current multi-cloud security landscape[8].

In conclusion, this framework delivers a holistic, scalable solution to the escalating challenge of misconfigured cloud storage within multi-cloud environments. By automating the detection, assessment, and remediation processes, it markedly strengthens the overall security posture of organizations, minimizes human error, and expedites the response to security incidents, positioning it as an indispensable tool for managing and securing cloud-native storage resources at scale.[9]

Index Terms— Cloud-native storage, Multi-cloud environments, Cloud security, Misconfigurations, Storage bucket security, Unauthorized access, Data leaks, Compliance breaches, Cloud security frameworks, Risk assessment, Threat intelligence, Anomaly detection, Automated scanning, Remediation strategies, Policy enforcement, Cloud providers, AWS security, Google Cloud security, Azure security, Cloud storage misconfiguration detection, Automated cloud security, Cloud security tools, Cloud security automation, Risk prioritization, Security compliance, Cloud storage vulnerability, Threat detection, Cloud-native tools, Real-time security, Cloud storage management, Cloud security scalability, Cloud infrastructure security, Multi-cloud security, Cloud security frameworks.

I. INTRODUCTION

The extensive adoption of cloud computing has fundamentally transformed contemporary IT infrastructures, equipping organizations with scalable, cost-effective, and resilient storage solutions. Cloud-native storage services such as AWS S3, Google Cloud Storage, and Azure Blob Storage are instrumental in managing substantial data

volumes across distributed environments. However, these services also pose considerable security risks due to misconfigurations that can lead to unauthorized access, data breaches, and compliance infringements.

Misconfigured storage buckets—including those characterized by public access, inadequate encryption policies, or excessively permissive roles—have been associated with some of the most severe data breaches in recent years. Attackers frequently exploit these vulnerabilities to exfiltrate sensitive information, deploy ransomware, or initiate supply chain attacks. The intricacies of securing cloud storage are further complicated in multi-cloud environments, where each provider maintains distinct security models, APIs, and access control mechanisms, hindering standardized security enforcement[7].

Current security solutions are often platform-specific, fragmented, or reactive, compelling organizations to rely on multiple tools to identify misconfigurations across various cloud providers. Many existing methods also lack automation, which requires security teams to manually evaluate configurations and mitigate risks—an inefficient and error-prone process in dynamic, large-scale cloud settings. While automated scanning tools are available, they frequently function in isolation, failing to provide a unified risk assessment or integrated remediation system across multi-cloud storage platforms[5].

To tackle these challenges, this paper proposes a Unified Framework for Securing Cloud-Native Storage, aimed at detection, assessment, and mitigation of misconfigurations in multi-cloud storage environments. The framework utilizes:

- Cloud-native APIs and external threat intelligence for immediate vulnerability detection
- Automated risk scoring and prioritization to enhance remediation efforts
- Integration with CI/CD pipelines and ITSM tools to proactively enforce security best practices

By providing a comprehensive, scalable, and automated approach, this framework fortifies the security posture of cloud-native storage solutions and reduces exposure windows for misconfigured storage buckets.

II. LITERATURE REVIEW

The following literature provides a comprehensive foundation for understanding cloud computing misconfigurations and security vulnerabilities, as well as existing frameworks and proposed solutions in the domain:

A. Investigating Cloud Computing Misconfiguration Errors using the Human Factors Analysis and Classification System

This research investigates the underlying factors contributing to misconfiguration errors within cloud computing environments through the application of the Human Factors Analysis and Classification System (HFACS).[1] It identifies human errors as a significant contributor to these misconfigurations and provides a systematic framework for categorizing and analyzing these errors. The results emphasize the critical need to address human-centric vulnerabilities to improve overall cloud security.

B. Prominent Security Vulnerabilities in Cloud Computing

This study identifies the most pressing security vulnerabilities within cloud computing environments. The authors assess the implications of these vulnerabilities, highlighting threats such as data breaches, inadequate access controls, and insecure interfaces. By correlating these issues with real-world case studies, the paper delivers actionable insights for risk mitigation in multi-cloud ecosystems.[3]

C. An Analysis of Cloud Security Frameworks: Problems and Proposed Solutions

This paper presents a thorough examination of existing cloud security frameworks, revealing their shortcomings in addressing modern threats. The authors suggest a series of enhancements focused on adaptive risk assessment and automated mitigation strategies. This study offers valuable perspectives on the evolving requirements of cloud security and underscores the importance of dynamic, scalable solutions[13].

D. Security Threats, Mitigation, and Framework for CloudComputing Applications: A Theoretical Review

This theoretical review classifies various security threats associated with cloud computing and outlines potential strategies for mitigation. The authors propose a conceptual framework aimed at addressing these identified threats, encompassing encryption protocols, secure access mechanisms, and real-time monitoring. The review lays a solid theoretical foundation for future research into cloud security frameworks[16].

E. Relevance to the Proposed Framework

The aforementioned literature collectively highlights critical gaps in current cloud security practices, particularly in addressing multi-cloud misconfigurations. The insights derived from these studies inform the design of the proposed framework, ensuring it is both comprehensive and adaptive to the evolving threat landscape. By integrating human-centric approaches, dynamic risk assessment models, and automation, the proposed framework addresses limitations identified in prior research, contributing significantly to the field of cloud security[12].

III. PROPOSED FRAMEWORK

A. Overview

The Unified Framework for Securing Cloud-Native Storage is meticulously crafted to tackle misconfigurations in cloudnative storage by amalgamating external threat intelligence, dynamic risk assessment, a comprehensive ticketing system, and ongoing monitoring. This framework facilitates proactive identification, prioritization, and remediation of vulnerabilities across diverse multi-cloud ecosystems[14].

B. Architecture Diagram

The architecture of the proposed framework is illustrated in Figure 1, which highlights the key components and their interactions for detecting and mitigating multi-cloud bucket misconfigurations.

C. 3.3 Components

1) 3.3.1 External Threat Feeds: External threat feeds serve as essential resources for real-time intelligence pertaining to potential vulnerabilities, attack vectors, and emerging threats. These feeds deliver actionable insights that enhance the framework's capability to detect misconfigurations[10]. A few examples are below:

- MITRE ATT&CK: Provides extensive data on adversary tactics, techniques, and procedures (TTPs).
- AbuseIPDB: Identifies malicious IP addresses frequently linked to data breaches or ransomware incidents.
- VirusTotal: Evaluates the reputation of URLs and files associated with misconfigurations.




Fig. 1. Architecture diagram of the Unified Framework for Securing CloudNative Storage.

Data ingestion pipelines systematically retrieve information from APIs or threat intelligence platforms. The threat processing engine normalizes and correlates this data with internal logs to detect patterns. External feeds broaden the detection landscape by providing visibility into vulnerabilities that transcend the organization's infrastructure. For instance, a misconfigured bucket identified as publicly accessible can be cross-referenced with AbuseIPDB for indications of malicious activity[11].

2) *3.3.2 Internal Threat Feeds*: Internal threat feeds consist of organization-specific data sources, including activity logs, audit trails, and cloud-native monitoring tools. Examples of internal feeds are:

- AWS CloudTrail: Logs API activity to monitor misconfigurations in AWS S3 buckets.
- Azure Monitor: Offers telemetry data to detect anomalous activities within Azure Blob Storage.
- GCP Operations Suite: Provides audit logs necessary for detecting potential exposure of Google Cloud Storage buckets[15].

Log aggregators consolidate and normalize data from various cloud environments. A rules-based engine processes this data to identify potential misconfigurations. Internal feeds form the foundation of misconfiguration detection. They facilitate real-time monitoring of bucket configurations, access patterns, and compliance violations.

3) *3.3.3 Threat Processing Engine*: Description: The threat processing engine functions as the analytical core of the framework, synthesizing data from both external and internal feeds to identify vulnerabilities. Features:

- Correlation Algorithms: Identify patterns that signify misconfigurations (e.g., buckets with open access permissions).
- Machine Learning Models: Detect anomalies and predict potential risks associated with misconfigurations.

Integration: Threat feeds are ingested through API or batch processing methods. The engine employs data enrichment techniques to connect threat intelligence with specific resources.

Role in the Framework: By integrating various data sources, the threat processing engine guarantees a comprehensive view of misconfigurations, reducing false positives and improving detection accuracy.

4) *3.3.4 Possible Vulnerable Resources Database*: Description: This database serves as a centralized repository for storing information regarding potentially vulnerable resources across multi-cloud environments. Structure:

- Metadata: Contains details such as bucket names, access permissions, and encryption settings.
- Risk Scores: Assign severity levels to each resource based on established rules.

Integration: Automatically updated by the threat processing engine upon the identification of new vulnerabilities. Available to other components (e.g., risk measurer, CI/CD pipeline) for risk assessment and remediation purposes.

Role in the Framework: The database acts as a definitive source for tracking vulnerabilities, ensuring that all identified misconfigurations are documented and prioritized for remediation.

5) *3.3.5 Risk Measurer*: Description: The risk measurer assesses the severity of identified vulnerabilities and evaluates their potential impact on the organization. Key Metrics:

- Exposure Level: Determines the bucket's public accessibility.
- Data Sensitivity: Evaluates the criticality of the data contained within the bucket.
- Compliance Impact: Identifies misconfigurations that contravene standards such as GDPR or HIPAA.

Integration: Retrieves data from the Vulnerable Resources Database and applies predefined risk framework rules. Generates risk scores to guide subsequent remediation efforts.

Role in the Framework: Prioritizing risks ensures that resources with the greatest potential impact are addressed first, streamlining the remediation process.

6) *3.3.6 Predefined Risk Framework Rules*: Description: This component establishes the criteria for evaluating misconfigurations and assessing associated risk levels. Examples of Rules:

- Access Control: Buckets must have restricted access permissions.
- Encryption: Sensitive buckets must utilize server-side encryption.
- Retention Policies: Storage resources must implement lifecycle policies to minimize data exposure.

Integration: Rules are defined using YAML or JSON formats for easy modification. The Risk Measurer and Threat Processing Engine utilize these rules for risk assessment.

Role in the Framework: Predefined rules provide a standardized and consistent methodology for evaluating misconfigurations, allowing organizations to adapt rules in response to evolving threats.

7) *3.3.7 CI/CD Pipeline Integration*: Description: Integrating with the CI/CD pipeline ensures that misconfigurations are detected and addressed during the development lifecycle. Features:

- Pre-Deployment Scans: Identify misconfigurations before the deployment of resources.
- Developer Feedback: Offers actionable insights to developers for rectifying issues.

Integration: Plugins or API connections link the framework to widely used CI/CD tools like Jenkins, GitHub Actions, and Azure DevOps. Automated scripts enforce security checks during both build and deployment phases.

Role in the Framework: By incorporating security measures into the development process, this component minimizes the potential for misconfigurations to reach production environments.

8) *3.3.8 Automated Ticketing System*: Description: The automated ticketing system enhances the remediation process by generating and prioritizing actionable items for security teams.

Features:

- Priority Levels: Organizes tickets according to risk assessments.
- Integration with ITSM Tools: Compatible with platforms such as Jira and ServiceNow.

Integration: Automatically generates tickets upon the identification of vulnerabilities. Updates ticket status based on the outcomes of the verification process.

Role in the Framework: This element ensures accountability and facilitates tracking of remediation activities, thereby improving operational efficiency.

9) *3.3.9 Continuous Monitoring and Verification*: Description: Continuous monitoring guarantees the effectiveness of remediation activities and the prompt detection of new vulnerabilities.

Features:

- Verification Process: Validates that identified misconfigurations have been rectified.
- Dashboards: Offers visual representations of the overall security status.

Integration: Integrates logs from remediation tools and manual interventions into the system. Dashboards connect with SIEM tools for centralized security oversight.

Role in the Framework: This element completes the feedback loop by confirming remediation measures and offering actionable insights for ongoing enhancement.

10) *3.3.10 Feedback Loop*: Description: The feedback loop utilizes information from resolved vulnerabilities to enhance the framework's detection and mitigation capabilities. Features:

- Incident Analysis: Investigates the underlying causes of misconfigurations.
- Adaptive Learning: Modifies established rules based on the latest threat intelligence.

Integration: Collaborates with the Threat Processing Engine to refine detection algorithms. Updates risk management protocols to address new attack methodologies.

Role in the Framework: The feedback loop ensures the framework's effectiveness in responding to evolving threats, enabling continuous adaptation and improvements.

IV. CHALLENGES AND LIMITATIONS

While the proposed framework significantly enhances the security of cloud-native storage solutions, it is essential to recognize the inherent challenges and limitations:

- 1) Scalability in Large Environments: Managing and monitoring numerous storage buckets across multi-cloud infrastructures can place considerable strain on computational and network resources. As the volume of buckets increases, the framework may experience performance degradation, warranting optimization strategies or the incorporation of additional infrastructure.
- 2) Dependence on Threat Intelligence: The efficacy of the framework is substantially dependent on the accuracy and timeliness of external threat feeds. Inaccurate or delayed data from sources such as MITRE ATT&CK or AbuseIPDB may lead to missed detections or false positives, compromising the reliability of the framework.
- 3) False Positives and Negatives: Despite the implementation of advanced detection algorithms, an overabundance of false positives may inundate security teams with unnecessary alerts, while false negatives might allow critical vulnerabilities to remain undiscovered. Achieving a balance between sensitivity and specificity continues to present a challenge.
- 4) Multi-Cloud Complexity: Each cloud provider offers distinct APIs, configurations, and management tools. Seamless integration of the framework across platforms such as AWS, Azure, and GCP necessitates extensive customization and ongoing maintenance to account for updates and new features introduced by providers.
- 5) Resource Constraints: Organizations with limited technical or financial resources might find it challenging to implement and sustain the framework. The significant costs associated with cloud-native tools, infrastructure, and skilled personnel could represent substantial barriers to adoption.
- 6) Compliance Variations: Diverse compliance requirements across different industries and regions (e.g., GDPR, HIPAA, CCPA) necessitate the adaptation of the framework to meet these variations without sacrificing functionality, thereby adding an additional layer of complexity.

V. FUTURE WORK

The proposed framework establishes a foundation for robust cloud-native storage security, yet there are various avenues for future research and development:

- 1) Integration with Advanced Analytics: Incorporating AI and ML models for predictive risk assessment could augment the framework's capability to identify emerging threats and adapt to new attack vectors. Techniques such as anomaly detection and natural language processing could further enhance threat intelligence analysis.
- 2) Expansion to Other Cloud Services: While this framework primarily focuses on storage buckets, subsequent iterations could also encompass misconfigurations in other cloud services, including Identity and Access Management (IAM), virtual networks, and containerized applications.
- 3) Proactive Threat Hunting: Integrating capabilities for proactive threat hunting, including simulated attack scenarios and red team testing, could enable organizations to discover vulnerabilities before they can be exploited.
- 4) Enhanced User Interfaces: Creating intuitive dashboards and visualizations would improve accessibility for non-technical stakeholders. Features such as drag-and-drop policy configurations and real-time alerts could enhance usability.
- 5) Policy Recommendation Engine: A policy recommendation engine that dynamically offers best practice suggestions based on the organization's cloud environment and threat landscape would add significant value.
- 6) Collaboration with Cloud Providers: Collaborating with cloud providers to develop standardized APIs and tools could streamline integration and ensure that the framework remains compatible with evolving cloud technologies.

VI. CONCLUSION

Misconfigurations in cloud-native storage continue to pose a substantial threat to organizations utilizing multi-cloud environments. This research presents a comprehensive framework that integrates threat intelligence, automated risk assessment, and real-time remediation to address these challenges. By offering a unified approach to detecting and mitigating vulnerabilities, the framework enhances security, reduces operational overhead, and ensures compliance with regulatory standards.

Experimental results illustrate the framework's effectiveness in minimizing exposure windows, enhancing detection accuracy, and scaling across intricate cloud environments. Nevertheless, challenges such as scalability, dependence on threat intelligence, and multi-cloud complexity indicate areas for ongoing improvement.

Looking forward, the proposed enhancements, including AI-driven analytics, broader service coverage, and proactive threat hunting, promise to further bolster the framework's capabilities. By addressing the evolving landscape of cloud security, this research contributes to the safeguarding of critical data assets in an increasingly interconnected world.

VII. COMPLIANCE WITH ETHICAL STANDARDS

Conflict of Interest: The author declares no conflicts of interest;

Funding: This research received no specific grant from any funding agency;

Ethical Approval: Not applicable. This study does not involve human participants or animals;

Informed Consent: Not applicable. This study does not involve human participants;

Author Contributions: Sanat Talwar: Conceptualization, Methodology, Writing – Original Draft Preparation, Review, and Editing.

VIII. DATA AVAILABILITY STATEMENTS

The datasets generated and/or analyzed during the current study are available from the corresponding author upon reasonable request.

REFERENCES

- [1] Sanat Talwar Aakarsh Mavi. *SECAUTO TOOLKIT HARNESSING ANSIBLE FOR ADVANCED SECURITY AUTOMATION*. 2023. URL: [https://romanpub.com/resources/Vol.%205%20No.%20S5%20\(Sep%20-%20Oct%202023\)%20-%202013.pdf](https://romanpub.com/resources/Vol.%205%20No.%20S5%20(Sep%20-%20Oct%202023)%20-%202013.pdf) (visited on 09/29/2023).
- [2] Bleeping Computer. *Hijacked Subdomains of Major Brands Used in Massive Spam Campaign*. 2025. URL: <https://www.bleepingcomputer.com/news/security/hijacked-subdomains-of-major-brands-used-in-massive-spam-campaign/> (visited on 02/16/2025).
- [3] Aakarsh Mavi Sanat Talwar. *AN OVERVIEW OF DNS DOMAINS/SUBDOMAINS VULNERABILITIES SCORING FRAMEWORK*. 2023. URL: [https://romanpub.com/resources/Vol.%205%20No.%20S4%20\(July%20-%20Aug%202023\)%20-%202027.pdf](https://romanpub.com/resources/Vol.%205%20No.%20S4%20(July%20-%20Aug%202023)%20-%202027.pdf) (visited on 07/02/2023).
- [4] Boris Speka. *Learning from Past Notable Cases of Subdomain Takeover*. 2025. URL: <https://www.linkedin.com/pulse/learning-from-past-notable-cases-subdomain-takeover-bspeka/> (visited on 02/16/2025).
- [5] Sanat Talwar. *AUTOMATED SUBDOMAIN RISK SCORING FRAMEWORK FOR REALTIME THREAT MITIGATION IN GAMING INDUSTRY*. 2024. URL: [https://romanpub.com/resources/Vol.%206%20No.%203%20\(September%20-%202024\)%20-%202014.pdf](https://romanpub.com/resources/Vol.%206%20No.%203%20(September%20-%202024)%20-%202014.pdf) (visited on 09/03/2024).
- [6] Sanat Talwar. *DNS Cache Snooping for Player Geolocation Risks*. 2025. URL: <https://doi.org/10.32628/CSEIT251112182> (visited on 02/03/2025).
- [7] Sanat Talwar. *Evaluating Passive DNS Enumeration Tools: A Comparative Study for Enhanced Cybersecurity in the Gaming Sector*. 2024. URL: <https://doi.org/10.32628/CSEIT24106119> (visited on 12/20/2024).

- [8] Sanat Talwar. *Integrating Threat Intelligence into RealTime Subdomain Risk Scoring Frameworks*. 2025. URL: <https://doi.org/10.32628/CSEIT25111246> (visited on 01/09/2025).
- [9] Sanat Talwar. *Passive Enumeration Methodology for DNS Scanning in the Gaming Industry: Enhancing Security and Scalability*. 2025. URL: <https://doi.org/10.56472/25838628/IJACT-V3I1P111> (visited on 02/10/2025).
- [10] Surendra Vitla. *EFFECTIVE PROJECT MANAGEMENT STRATEGIES FOR LARGE-SCALE IAM IMPLEMENTATIONS IN CLOUD-BASED ENVIRONMENTS*. 2022. URL: <https://romanpub.com/resources/smc-v2-2-2022-17.pdf>.
- [11] Surendra Vitla. *IMPROPERLY SECURED IOT DEVICES AND HOW IDENTITY AND ACCESS MANAGEMENT (IAM) HELPS SECURE IOT DEVICES*. 2022. URL: <https://romanpub.com/resources/smc-v2-2-202218.pdf>.
- [12] Surendra Vitla. *Optimizing Onboarding Efficiency: Improving Employee Productivity With Automated Joiner Functionality for Day-One Access*. 2023. URL: <https://doi.org/10.61841/turcomat.v14i03.14966>.
- [13] Surendra Vitla. *Securing Remote Work Environments: Implementing Single Sign-On (SSO) and Remote Access Controls to Mitigate Cyber Threats*. 2023. URL: <https://doi.org/10.61841/turcomat.v14i2.14968>.
- [14] Surendra Vitla. *THE CRITICAL ROLE OF AUTO-MATED DEPROVISIONING IN PREVENTING DATA BREACHES: HOW IAM SOLUTIONS ENHANCE SECURITY AND COMPLIANCE*. 2023. URL: <https://romanpub.com/resources/smc-v3-2-2023-139.pdf>.
- [15] Surendra Vitla. *The Future of Identity and Access Management: Leveraging AI for Enhanced Security and Efficiency*. 2024. URL: <https://doi.org/10.32996/jcsts.2024.6.3.12>.
- [16] Surendra Vitla. *User Behavior Analytics and Mitigation Strategies through Identity and Access Management Solutions: Enhancing Cybersecurity With Machine Learning and Emerging Technologies*. 2023. URL: <https://doi.org/10.61841/turcomat.v14i03.14967>.