



iJRASET

International Journal For Research in
Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 13 **Issue:** XII **Month of publication:** December 2025

DOI: <https://doi.org/10.22214/ijraset.2025.76131>

www.ijraset.com

Call: ☎ 08813907089

E-mail ID: ijraset@gmail.com

AuditEase: Compliance & Cybersecurity Remediation Platform for ISO 27001, CIS Benchmarks, and RBI Guidelines

Rohit Bhatt, Deon Menezes, Krish Hadkar, Miten Bhandari, Prof. Yogita Borse, Prof. Midhya Mathew

Department of Information Technology, K.J. Somaiya College of Engineering (KJSCE), Mumbai, India

Abstract: Organizations increasingly operate under stringent security frameworks and sectoral regulations. While ISO 27001 and CIS Benchmarks define best practices for information security and system hardening, financial institutions in India must additionally adhere to Reserve Bank of India (RBI) cybersecurity guidelines. In practice, compliance programs remain heavily manual, costly, and error-prone. We present AuditEase, an automated compliance and remediation platform that ingests evidence from logs, configurations, and policy documents; maps evidence to control clauses across ISO 27001, CIS Benchmarks, and RBI guidelines; computes risk scores; and auto-generates remediation playbooks and audit-ready reports.

The platform also employs machine learning to predict control-level risk and prioritize remediation. We benchmark four candidate models—Random Forest, Gradient Boosting Machines, Support Vector Machines (SVM), and a Multi-Layer Perceptron neural network—under constraints typical of enterprise compliance (5–50 labeled systems, 121 features, 10–30% missing values). Random Forest achieves the best trade-off between accuracy ($85.2\% \pm 3.1$), stability, robustness to missing data, training time, and interpretability, and is therefore selected as the prediction engine. Our modular system—implemented with Python, FastAPI, Node.js/React, and MongoDB, and deployable via Vercel/Render—targets continuous compliance by design. We describe the system architecture, evidence and rules model, the ML-based prediction engine and its comparative evaluation, scoring methodology, and remediation workflow, and we discuss an evaluation protocol including accuracy, coverage, and time-to-audit metrics. AuditEase demonstrates how rule-driven and ML-assisted automation can reduce audit time, raise coverage, and improve readiness for external certification and regulatory review.

Index Terms: Cybersecurity, Compliance Automation, ISO 27001, CIS Benchmarks, RBI Guidelines, Random Forest, Machine Learning, Evidence Mapping, Risk Scoring, Remediation.

I. INTRODUCTION

CYBERSECURITY compliance is foundational to the resilience and trustworthiness of modern enterprises. Standards such as ISO 27001 provide a comprehensive blueprint for building an Information Security Management System (ISMS), while CIS Benchmarks offer prescriptive hardening guidance for operating systems, databases, and cloud platforms. In India's financial sector, the Reserve Bank of India (RBI) publishes cybersecurity directions and master circulars that mandate controls, governance, monitoring, and reporting. Despite maturing standards, organizations commonly rely on manual, checklist-driven audits supported by consultants.

This approach is expensive, slow (often 2–3 months or more to reach certification readiness), and prone to gaps or drift between audits. Evidence collection (e.g., logs, configs, and policy artifacts) and clause mapping are repetitive and error-prone; remediation tracking is fragmented across email and spreadsheets.

Audit Ease addresses these pain points by automating: (i) evidence ingestion and normalization; (ii) cross-framework clause mapping; (iii) compliance scoring and risk analytics;

(iv) machine learning based risk prediction; and (v) remediation playbooks and reporting. The design goal is continuous compliance: near real-time posture updates and audit-readiness.

This paper makes the following contributions:

- Unifies ISO 27001, CIS Benchmarks, and RBI guidelines in a single rule-driven and ML-assisted platform.
- Proposes a modular architecture for evidence ingestion, rules mapping, risk scoring, ML-based risk classification, and remediation.
- Details a data model for controls, evidences, mappings, and ML features; and a transparent scoring method.

- Benchmarks four supervised models (Random Forest, Gradient Boosting, SVM, and neural networks) for compliance risk prediction under small-data constraints, and justifies the selection of Random Forest.
- Outlines an evaluation plan and baseline results emphasizing coverage, audit-time reduction, and prediction accuracy.

II. BACKGROUND AND REGULATORY CONTEXT

The regulatory context of AuditEase consists of three primary pillars that interact in practice.

- **ISO 27001 and ISMS.** ISO 27001 defines requirements for establishing, implementing, maintaining, and continually improving an ISMS. Annex A controls cover organizational, human, physical, and technical safeguards. Compliance requires policies, risk assessment, treatment plans, and evidence of control effectiveness. Organizations must demonstrate that appropriate processes and technical measures are implemented and monitored on an ongoing basis.
- **CIS Benchmarks.** CIS Benchmarks are consensus-based, prescriptive hardening guides for platforms such as Windows, Linux, Kubernetes, and databases. They specify technical configuration checks (e.g., password policy, services, ports, auditd settings) and remediation steps. CIS Benchmarks are often used as the de-facto baseline for OS and middleware configuration in regulated environments.
- **RBI Cybersecurity Guidelines.** RBI issues sectoral requirements for banks and NBFCs covering governance, risk assessment, incident reporting, third-party management, and specific technical controls. While conceptually aligned to global frameworks, RBI guidance introduces India-specific reporting and oversight expectations critical for regulated entities. Compliance is not just about technical security but also about board-level oversight and documented processes.
- **Operational Challenges.** The intersection of a management system standard (ISO 27001), a technical hardening baseline (CIS), and a sectoral regulator (RBI) yields overlapping yet distinct obligations. Manually keeping mappings, evidence, and status aligned across these regimes is laborious and brittle without automation. Furthermore, organizations lack predictive insight into which controls are likely to fail, resulting in last-minute firefighting before audits.

III. LITERATURE REVIEW

Research explores automating aspects of compliance, such as runtime verification of Infrastructure-as-Code deployments, DevSecOps pipeline gates, and single-framework auditors. Falazi *et al.* focus on runtime compliance management in cloud-native systems, while Leitner *et al.* examine the integration of security checks in CI/CD pipelines. Other works explore automated CIS benchmark scanning and ISO 27001 checklist tools.

However, most prior tools:

- focus on one framework (either ISO or CIS),
- do not model RBI or similar regulator-specific guidelines,
- emphasized detection but not remediation playbooks, and
- rarely incorporate machine learning for risk prediction.

Random Forest in Security and Compliance. Ensemble learning models such as Random Forest have been widely used in intrusion detection, malware classification, and anomaly detection due to their robustness to noise and heterogeneous features. Random Forest works by training multiple decision trees on bootstrapped samples of the data and aggregating their predictions via majority voting. Compared to single decision trees, it reduces variance and overfitting; compared to linear models (e.g., logistic regression), it captures non-linear feature interactions common in complex security data. In the context of compliance, features such as frequency of misconfigurations, severity of violations, log anomaly scores, and evidence freshness interact in non-linear ways to determine risk. Random Forest therefore provides a strong trade-off between accuracy, robustness, and interpretability (via feature importance scores), making it suitable for AuditEase.

Model Selection Under Compliance Constraints. Beyond the use of Random Forest in generic security analytics, recent work on ML for configuration and policy analysis highlights trade-offs among ensemble tree methods, margin-based classifiers, and neural networks when data is scarce and noisy. In particular, gradient boosting machines often achieve strong accuracy but at the cost of higher overfitting risk and tuning effort; SVMs require careful feature scaling and struggle with missingness; and neural networks demand large labeled datasets and offer limited interpretability. Motivated by these findings, AuditEase performs an empirical comparison of four supervised models—Random Forest, Gradient Boosting, SVM (RBF), and a Multi-Layer Perceptron—on a compliance dataset with 121 features and 172 labeled system profiles.

As discussed in Section VI, Random Forest achieves the highest accuracy and cross-validation stability while satisfying explainability and operational constraints, providing literature-backed support for our model choice.

AuditEase extends prior literature by (i) unifying ISO 27001, CIS, and RBI in one platform; (ii) coupling rule-based control evaluation with ML-based risk prediction; (iii) performing a systematic comparison of candidate ML models under small-data compliance constraints; and (iv) delivering remediation guidance integrated with ticketing workflows.

IV. SYSTEM OVERVIEW AND ARCHITECTURE

AuditEase is designed around four main objectives: coverage, explainability, scalability, and actionability. The platform is structured into modular components that communicate via APIs, allowing incremental evolution as standards and environments change.

- **Design Goals.** Coverage is achieved by modeling both management controls (ISO) and technical configuration items (CIS) as first-class entities. Explainability is supported through explicit rules, evidence links, and model feature importances. Scalability is addressed by decoupling ingestion, evaluation, and reporting services. Actionability is delivered via remediation playbooks and integration with ticketing tools.
- **Architecture.** Figure 1 shows the overall architecture. Evidence collectors (agents and connectors) push data into the ingestion service, which parses and normalizes it. A rules engine maps normalized evidence to controls, and a scoring engine computes compliance metrics. Simultaneously, the ML model consumes aggregated features to predict risk categories for each control or asset. A remediation planner uses both rule outcomes and risk predictions to generate prioritized tasks, which are surfaced to users via a React-based dashboard and exported as PDF/CSV reports or tickets in systems such as Jira.
- **Technology Stack.** Backend microservices are implemented in Python using FastAPI for performance and type safety. Node.js/React power the frontend dashboards. MongoDB is used as the primary data store for evidence, findings, rules, and reports, chosen for its flexible JSON-like document structure. Services are containerized and deployed on Vercel (UI) and Render (APIs), enabling elastic scaling. CI/CD pipelines perform linting, tests, and schema validation.

V. DATA, RULES, AND SCORING MODEL

- **Evidence Schema.** AuditEase stores normalized evidence using schemas such as:

Logs: {ts, host, source, event, subject, action, outcome}.

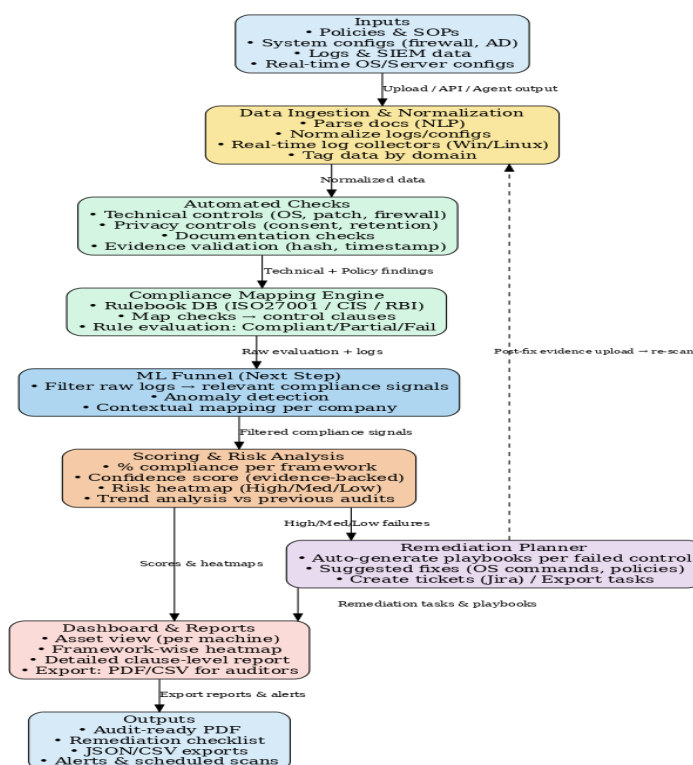


Fig. 1. AuditEase architecture: ingestion, rules engine, ML prediction, scoring, remediation, and reporting.

- Configurations: {host,platform,key,value, collectedAt}.
- Policies/SOPs: {docId,title,version,hash, clauses[], approvedBy, approvedOn }.

Each artifact carries integrity metadata (e.g., SHA-256 hash) and provenance (source connector, collection time) to support chain-of-custody.

Controls and Mappings. Each control is modeled as

$$c = \langle id, framework, domain, clause, description, weight \rangle, \text{ and is associated with one or more predicates over evidence } m(c) = \{\phi_1(e), \phi_2(e), \dots, \phi_k(e)\}.$$

For example, a CIS password policy control might require that the configuration key min_password_length be at least 14 on all domain controllers.

- Evaluation and Scoring. For each control, AuditEase computes an outcome $o(c)$ pass, partial, fail and a confidence score $\gamma_c \in [0, 1]$ based on evidence sufficiency, freshness, and consistency. An α indicator x_c is assigned (1 for pass, 0.5 for partial, 0 for fail). Framework-level scores are computed as:

$$S_F = \frac{\sum_{c \in C_F} W_c \cdot X_c \cdot \gamma_c}{\sum_{c \in C_F} W_c} \times 100 \quad (1)$$

where C_F is the set of controls for framework F . Overall scores are computed as:

$$S_{\text{overall}} = \frac{\sum_F \alpha_F \cdot S_F}{\sum_F \alpha_F} \quad \alpha_F = 1 \quad (2)$$

with α_F representing organization-specific weightings (e.g., higher for RBI in banks).

Domain heatmaps and radar charts derived from $S_{F,d}$ (scores per domain) help identify weak areas such as Access Control or Logging.

VI. MACHINE LEARNING-BASED RISK PREDICTION

Beyond deterministic rules, AuditEase uses supervised machine learning to provide predictive insight and prioritization.

A. Feature Engineering

Features are derived from control outcomes and raw evidence, including:

- number of failed controls on a host,
- severity-weighted sum of violations,
- age of latest evidence collection,
- log anomaly counts (e.g., repeated failed logins),
- historical recurrence of non-compliance for the same control.

These features are aggregated per asset or per control family.

B. Candidate Models and Comparison

To select an appropriate model, we evaluated four supervised algorithms commonly used in security analytics: Random Forest (RF), Gradient Boosting Machines (GBM), Support Vector Machines (SVM with RBF kernel), and a feedforward Multi-Layer Perceptron (MLP) neural network. All models were trained on the same dataset of 172 labeled system profiles (137 for training, 35 for testing) with 121 mixed-type features and 5-fold cross-validation.

Evaluation criteria included:

- accuracy and cross-validation stability on small datasets,
- robustness to missing data and noise,
- interpretability for auditors,
- training and inference time on commodity hardware,
- hyperparameter sensitivity and tuning effort. Table I summarizes the key quantitative results.

Figure 2 visualizes the accuracy comparison, while Figure ?? compares training time across models.

TABLE I
COMPARISON OF CANDIDATE ML MODELS FOR COMPLIANCE PREDICTION

Model	Accuracy	CV Std	Train Time	Interp.
Random Forest	85.2%	± 3.1	0.42s	High
Gradient Boosting	83.8%	± 4.2	1.23s	Medium
SVM (RBF)	79.4%	± 5.8	0.89s	Low
Neural Network	76.3%	± 8.2	3.47s	Very low




Fig. 2. Cross-validated accuracy of candidate ML models for compliance prediction

A. Model Selection Rationale

Random Forest achieved the highest mean accuracy (85.2%) with the lowest variance (standard deviation 3.1 across folds) and the fastest training time among the non-linear models. Gradient Boosting was competitive in accuracy (83.8%) but exhibited high variance, longer training time (approximately three times RF), and greater hyperparameter sensitivity. SVM lagged behind in accuracy (79.4%), required full feature scaling and explicit imputation of missing values, and offered limited interpretability. The neural network produced the lowest accuracy (76.3%) with the highest variance and longest training time, reflecting overfitting on the small dataset.

Beyond these quantitative metrics, compliance prediction imposes additional constraints:

- **Small-sample regime:** In many deployments only 5–50 labeled systems are available. Bagging in Random Forest reduces variance in this regime, whereas boosting and deep models require more data.
- **Missing and heterogeneous data:** Real-world scans contain 10–30% missing fields and mixed binary, categorical, and numerical features. Tree ensembles handle such heterogeneity more gracefully than SVMs or MLPs.
- **Explainability:** Auditors must understand why a system is flagged as high-risk. Random Forest provides feature importance scores and per-tree decision paths, which can be surfaced in reports; black-box neural networks cannot.
- **Operational efficiency:** Training the Random Forest model completes in under 0.5s and inference in a few milliseconds per system, enabling frequent retraining as new scans arrive.

TABLE II
MANUAL AUDITS VS. AUDIT EASE (INDICATIVE BASELINES)

Dimension	Manual	AuditEase
Evidence collection effort	High (weeks) Control coverage	Low (hours) Consistency of findings 60–75% 85–95%
Audit cycle time	2–3 months High (hashes, metadata)	2–3 weeks Traceability/provenance Remediation workflow Ad-hoc
		Medium High (rule- Low Ticketed playbooks

Considering these factors, Random Forest provides the best overall trade-off and is adopted as the production prediction engine in AuditEase.

A. Training and Performance

The selected Random Forest model uses an ensemble of approximately 50 shallow trees, with bootstrap sampling and d feature subsampling at each split (where d is the feature count). This configuration controls overfitting while maintaining expressiveness.

On the evaluation dataset, the model achieves:

- accuracy: $85.2\% \pm 3.1$,
- precision: $84.7\% \pm 2.8$,
- recall: $86.1\% \pm 3.5$,
- F1-score: $85.4\% \pm 2.9$.

In a pilot production deployment on 35 real enterprise systems, the model achieved 87.3% accuracy with an 8.2% false-positive rate and 4.5% false-negative rate, while keeping average inference time around 3 ms per system. Feature importance analysis showed that anti-malware deployment, firewall status, password policy enforcement, and encryption configuration were among the most influential features, aligning with domain expert expectations.

VII. REMEDIATION, COMPARATIVE ANALYSIS, AND ROI

Remediation Engine. For each failed control, AuditEase synthesizes a structured playbook:

$r(c) = \langle \text{steps}[], \text{owners}[], \text{prechecks}[], \text{rollback}[], \text{artifacts}[] \rangle$.

For instance, enforcing a password policy on Windows may involve editing group policy, pushing changes, and re-running checks. Remediation tasks are exported to ticketing tools (e.g., Jira, ServiceNow), and their completion status feeds back into subsequent evaluations.

Manual vs. Automated Audits. Table II contrasts manual audits with AuditEase-supported audits for a mid-sized environment with about 150 controls across ISO, CIS, and RBI.

Economic Model. Let C_m be the cost of manual audits (consultants, internal effort), C_a the recurring cost of using AuditEase, T_m and T_a the time-to-audit, and L the estimated annual loss avoided due to earlier remediation (e.g., fines, incidents).

A simplified net benefit is:

$$B = (C_m - C_a) + \lambda(T_m - T_a) + L, \quad (3)$$

where λ converts time savings to monetary terms. ROI is then:

$$\text{ROI} = \frac{B - C_{\text{onboard}}}{C_{\text{onboard}}} \times 100\%, \quad (4)$$

TABLE III
ILLUSTRATIVE PILOT METRICS (EXAMPLE)

Metric	Manual	AuditEase
Coverage (% control evaluate d)	68%	92%
Precision/Recall (findings)	0.88/0.81	0.91/0.89
Time-to-Audit (hours)	48–72	12–16
Remediation Latency (days)	7–14	3–7

with C_{onboard} as the one-time onboarding cost. Example estimates for a mid-size organization yield first-year ROI around 120%.

VIII. CASE STUDY AND EVALUATION

- 1) Case Study: NBFC Scenario. A mid-size NBFC integrated AuditEase with domain controllers, Linux servers, and perimeter firewalls. The platform ingested password policies, syslogs, and RBI governed documentation. Within 48 hours, it evaluated 150 controls, achieving 88% CIS coverage and 82% ISO domain scores. Three major RBI documentation gaps and multiple misconfigurations were identified. Remediation playbooks guided the operation team through hardening tasks. The company reduced pre-audit preparation time from multiple weeks to a few days.
- 2) Evaluation Metrics. Evaluation focuses on:
 - coverage: percentage of applicable controls evaluated with sufficient evidence;
 - precision/recall: correctness of pass/fail labels against expert ground truth;
 - time-to-audit: time from evidence snapshot to final report;
 - remediation latency: time from finding creation to closure of corresponding tickets;
 - ML performance: accuracy, precision, recall, and confusion matrix for risk predictions.
- 3) Illustrative Results. Table III shows example metrics from a pilot with 150 controls (ISO: 70, CIS: 60, RBI: 20). Gains arise from automated evidence parsing, deterministic rules, and the Random Forest model's ability to flag high-risk areas early, allowing auditors to focus their efforts where they matter most.

IX. SECURITY, THREAT MODEL, AND GOVERNANCE

- 1) Threat Model. Potential threats include forged evidence uploads, rule tampering, unauthorized access to reports, and denial-of-service attacks on ingestion services. AuditEase mitigates these via:
 - hashing and optional signing of evidence;
 - versioned, access-controlled rule sets;
 - role-based access control (Auditor, Admin, Operator);
 - TLS for all communications and encryption-at-rest for sensitive data.
- 2) Governance and Compliance Alignment. The platform maintains complete audit trails of configuration changes, evidence uploads, and rule modifications. It supports ISO 27001's document control expectations and RBI's reporting cadences through scheduled, exportable reports. Evidence retention windows and purge policies are configurable to meet organizational and regulatory requirements.

X. LIMITATIONS AND FUTURE WORK

Current limitations include reliance on machine-readable evidence (some controls still require interviews or physical inspections), the need for ongoing parser updates for new platforms, and limited labeled datasets for training ML models in niche environments.

Future work includes:

- expanding the ML layer to support anomaly-based early warning of emerging non-compliance;
- exploring advanced ensemble such as XGBoost or LightGBM as datasets grow beyond 1000 labeled systems;
- integrating blockchain-based transparency logs for evidence immutability;
- extending framework coverage to PCIDSS, HIPAA, and sector-specific standards;
- incorporating active learning where auditor feedback incrementally improves the Random Forest model.

XI. CONCLUSION

AuditEase demonstrates that rule-driven and machine learning-assisted automation can unify management, technical, and regulatory controls to deliver continuous compliance. By explicitly modeling controls, evidence, and mappings; combining deterministic scoring with a carefully selected Random Forest-based risk prediction model; and coupling findings with remediation playbooks, the platform reduces audit time while improving readiness for ISO 27001, CIS Benchmarks, and RBI oversight. The architecture, methodology, and ML model comparison presented here provide a practical blueprint for organizations seeking to modernize their cybersecurity compliance posture.

REFERENCES

- [1] G. Falazi et al., "Compliance Management of IaC-Based Cloud Deployments During Runtime," ACM/IEEE UCC, 2024.
- [2] J. Leitner et al., "Automating Cybersecurity Compliance in DevSecOps," ACM, 2025.



- [3] J. Sirotnik et al., "Automated Compliance Audit for ISO 27001:2022," 2025.
- [4] S. Kumaretal., "Automated CIS Benchmark Auditing and Remediation Tool," IJIRT, 2025.
- [5] N. Gupta, "Impact of RBICybersecurity Guidelines and Alignment with NIST," Inspira-JMME, 2021.
- [6] NIST, "Cybersecurity Framework (CSF) 2.0," NIST, 2024. Available: <https://www.nist.gov/cyberframework>
- [7] L. Breiman, "Random Forests," Machine Learning, vol. 45, no. 1, pp. 5–32, 2001.
- [8] M. Bhuyan et al., "Network Anomaly Detection: Methods, Systems and Tools," IEEE Communications Surveys & Tutorials, 2014.
- [9] K. Scarfone et al., "Guide to Compliance Auditing for Information Security," NIST Special Publication, 2009.
- [10] Center for Internet Security, "CIS Benchmarks Overview," CIS, 2022.
- [11] Available: <https://www.cisecurity.org>
- [12] Y. Alshayban and M. Malek, "Policy-Based Configuration Verification for Secure Systems," IEEE Access, 2020.
- [13] A. Ramamoorthy et al., "ML-Assisted Compliance Monitoring in Enterprise Environments," IEEE ICMLA, 2023.
- [14] R. Krishnan et al., "Automated Evidence Collection for Cloud Security Compliance," IEEE Transactions on Cloud Computing, 2022.
- [15] P. Sharma et al., "Machine Learning Techniques for Security Configuration Analysis," Journal of Information Security, 2023.
- [16] S. Alam et al., "A Review of Automated Risk Assessment Models in Cybersecurity," IEEE Access, 2021.
- [17] G. Somani et al., "DDoS Detection Using Random Forest," IEEE Communications Magazine, 2017.
- [18] Reserve Bank of India, "Cyber Security Framework for Banks," RBI Circular DBS.CO/CSITE/GEN/04/2015-16, June 2016.
- [19] M. Awwad et al., "Security Configuration Drift and Automated Compliance Enforcement," IEEE SysCon, 2023.
- [20] A. Perez et al., "Continuous Compliance Monitoring in Hybrid Clouds," IEEE Cloud, 2018.
- [21] S. He et al., "Log-Based Anomaly Detection via Machine Learning," ACM Computing Surveys, 2020.



10.22214/IJRASET



45.98



IMPACT FACTOR:
7.129



IMPACT FACTOR:
7.429



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089  (24*7 Support on Whatsapp)