



# Detection Decay in Multi-Tenant SIEM: Measuring Rule Half-Life and a Decay-Aware Content Lifecycle Model

Suman Sah

Cybersecurity Architect, Ernst & Young, Dallas, TX, USA.

Received On: 02/05/2025

Revised On: 12/05/2025

Accepted On: 27/05/2025

Published on: 17/06/2025

**Abstract** - Enterprise SIEM detection content is authored once and audited rarely, yet its efficacy degrades continuously as log schemas drift, business processes change, and adversaries adapt. Existing detection-engineering practice treats coverage as a static ATT&CK matrix property and schedules rule review on fixed calendars, implicitly assuming efficacy is time-invariant. This paper introduces detection decay as a measurable property of deployed correlation logic. We define rule half-life as the elapsed operational time at which a detection's true-positive yield falls to half its post-tuning baseline, and estimate it using Cox proportional-hazards survival analysis with time-varying covariates over 24 months of alert-disposition telemetry from a multi-tenant managed-SIEM estate (11 tenants, 1,847 deployed detections, 2.1M dispositions) spanning Splunk Enterprise Security and Microsoft Sentinel. Covariates - schema-change events, ingestion-volume shift, field-conformance failure rate, and tuning interventions - are extracted automatically from CIM/ASIM conformance telemetry and index-time parsing errors, yielding a per-rule hazard score computable inside the SIEM itself. Results show decay is neither uniform nor primarily adversary-driven: schema and telemetry-source instability account for 61% of hazard in credential-access and persistence detections, whereas defence-evasion logic decays fastest under adversary pressure (median half-life 94 days versus term-wide 211). We then propose a decay-aware lifecycle governance model that replaces calendar-based review with hazard-triggered revalidation under a bounded budget. Against a fixed-interval baseline it retains 87.4% detection efficacy versus 68.1% ( $p < 0.001$ ), reduces analyst-hours per true positive by 34.2%, and lowers false-positive volume by 30.6% at equal reviewer headcount.

**Keywords** - Detection Engineering, SIEM, Survival Analysis, Rule Half-Life, MITRE ATT&CK, Schema Drift, Security Operations.

## 1. Introduction

### 1.1. Background

Managed security service providers operate shared detection libraries across heterogeneous tenants. A single correlation rule may execute against dozens of log sources whose parsers, field names, and volumes evolve independently of the rule author. Detection content is therefore a deployed software artifact in a continuously mutating input environment. Security operations centres measure this artifact at a point in time - precision at tuning, ATT&CK coverage at audit - and then assume stability. Longitudinal measurement of deployed detection efficacy is largely absent from both practice and literature.

### 1.2. Limitations of Existing Approaches

Coverage-matrix methods score a rule set by the ATT&CK techniques it nominally touches, treating presence as efficacy. Detection-as-code pipelines apply CI/CD rigour to authoring but validate against synthetic or replayed data, not production drift. Purple-team exercises sample efficacy episodically at high cost and low frequency. Alert-triage machine learning suppresses false positives downstream without diagnosing upstream rule degradation. Emerging alternatives - Sigma-based portable rules, ASIM/CIM normalisation layers, and adversary-emulation platforms such as Atomic Red Team - reduce drift exposure but provide no

estimator of remaining rule lifetime and no mechanism to prioritise finite revalidation capacity.

### 1.3. Proposed Contribution

We reformulate detection efficacy as a time-to-event problem. Each rule is a subject; failure is the epoch at which its true-positive yield falls below half its post-tuning baseline. Fitting a Cox model with time-varying covariates produces a per-rule hazard score and an estimated half-life, decomposable into schema-induced and adversary-induced components. This decomposition is novel: it separates decay the provider can prevent from decay the provider must chase. We then convert hazard into a scheduling policy - a knapsack-constrained revalidation queue - and demonstrate that hazard-ranked review dominates calendar review at identical analyst cost.

## 2. Related Work and Background

### 2.1. Conventional Approaches

Traditional SIEM governance derives from log-management compliance. Rules are versioned in ticketing systems, reviewed annually or quarterly, and assessed by false-positive count. Efficacy is inferred from analyst complaint volume rather than measured. Coverage is reported as a technique count against ATT&CK, a metric insensitive to whether the underlying rule still fires. Threshold tuning is

reactive: a rule is examined when it floods the queue, not when it silently stops firing. Silent failure - the dominant decay mode - is structurally invisible to this regime.

## 2.2. Newer and Modern Approaches

Detection-as-code introduces version control, unit tests, and automated deployment for Sigma or KQL content. Normalisation schemas (Splunk CIM, Microsoft ASIM, OCSF) decouple rule logic from source-specific field names. Continuous validation platforms replay adversary emulation against production pipelines to confirm firing. These advances materially reduce drift but remain binary and episodic: a rule either passed the last test run or did not. None model the rate at which passing rules approach failure, and none allocate scarce revalidation capacity by expected loss.

## 2.3. Related Hybrid and Alternative Models

bility-remediation timing, patch-decay studies, and software defect prediction. Reliability-centred maintenance schedules inspection by failure hazard rather than calendar. Concept-drift detection in streaming ML monitors distributional shift in model inputs. Each is directly analogous to the detection problem, yet none has been transferred to SIEM content. The nearest prior work applies drift detection to ML-based detectors, not to human-authored correlation logic, and does not exploit normalisation-layer conformance telemetry as a drift signal.

## 2.4. Summary of Research Gap

Three gaps persist. First, detection coverage is modelled statically; no published estimator of rule lifetime exists. Second, decay causation is undecomposed - practitioners cannot distinguish a rule broken by a parser change from one evaded by an adversary, though the remedies differ entirely. Third, revalidation scheduling is unoptimised, allocating equal review effort to rules with radically different hazard. A multi-tenant MSSP vantage point uniquely permits cross-organisational estimation, isolating tenant-specific from content-intrinsic decay.

# 3. Proposed Methodology

## 3.1. Formal Definition of Rule Half-Life

Let rule  $i$  have post-tuning baseline true-positive yield  $y_i(0)$ , measured over a 30-day stabilisation window following its last validated tuning event. Yield is defined as confirmed true positives per 1,000 executions, normalised by tenant ingestion volume. Failure occurs at the first epoch  $t$  where the 14-day trailing mean of  $y_i(t)$  falls below  $0.5 \cdot y_i(0)$  and remains below for two consecutive windows. Half-life  $t_{1/2}^{(i)}$  is that epoch.

## 3.2. Hazard Model Specification

We fit a Cox proportional-hazards model  $h_i(t) = h_0(t) \exp(\beta_{\text{top}} x_i(t))$ , where  $x_i(t)$  is a time-

varying covariate vector. Covariates comprise schema-change count, CIM/ASIM field-conformance failure rate, index-time parsing-error rate, ingestion-volume z-shift, tenant vertical, ATT&CK tactic, data-source class, and cumulative tuning interventions. Time-varying specification is essential: schema changes are discrete shocks, not baseline attributes. Rules retired or superseded before failure are right-censored.

## 3.3. Decay Attribution

Hazard is decomposed by covariate group. Schema-attributable hazard is the exponentiated contribution of conformance, parsing-error, and volume covariates; residual hazard after conditioning on these is treated as adversary-attributable. Attribution is reported per tactic, enabling engineering versus threat-research triage.

## 3.4. Decay-Aware Governance Policy

Each rule receives hazard score  $\lambda_i = \exp(\beta_{\text{top}} x_i(t))$ , recomputed nightly. Revalidation is scheduled as a bounded knapsack: maximise expected retained efficacy  $\sum \lambda_i \cdot v_i$  subject to  $\sum c_i \leq B$ , where  $v_i$  is tenant-weighted rule value,  $c_i$  estimated engineer-hours, and  $B$  the sprint budget. This replaces fixed 90-day review.

## 3.5. Evaluation Design

A 24-month retrospective cohort is split into a 12-month fitting period and a 12-month prospective evaluation. The baseline arm applies fixed-interval review; the treatment arm applies hazard-ranked scheduling at identical budget. Primary endpoints are retained efficacy and analyst-hours per true positive; log-rank and paired  $t$ -tests establish significance.

The pipeline is deliberately closed-loop. Redeployment resets the stabilisation window, so every remediated rule re-enters the survival cohort as a fresh subject. This prevents survivorship bias from long-lived rules that have been silently repaired.

Covariate extraction is the architecturally novel stage. Conformance and parsing telemetry already exist in both platforms as operational health data and are conventionally discarded after dashboards render them. Repurposing them as hazard covariates requires no new instrumentation, which is what makes the method deployable inside an existing MSSP estate.

The fork after attribution encodes the operational payoff. Schema-attributable decay routes to platform engineering, where remedies are cheap and generalise across tenants. Adversary-attributable decay routes to threat research, where remedies are expensive and tenant-specific. Misrouting is the principal cost of undecomposed decay measurement.




Fig 1: Detection Decay Measurement and Governance Pipeline

## 4. Technical Implementation

### 4.1. Telemetry Layer

Disposition events are exported from Splunk ES notable-event review status and Sentinel incident classification into a unified schema keyed on rule\_uid, tenant\_id, epoch\_day, executions, alerts, tp\_confirmed. Field-conformance failures are derived from CIM datamodel acceleration audit logs and ASIM parser validation output. Parsing errors are read from index-time\_internal splunkd events and Sentinel ingestion-failure metrics.

### 4.2. Panel Construction

A daily panel is materialised per rule-tenant pair, producing counting-process rows (start, stop, event, covariates) suitable for time-varying Cox estimation. Volume z-shift is computed against a trailing 90-day tenant baseline to avoid conflating seasonal business cycles with drift.

### 4.3. Estimation Service

The model is fitted weekly using lifelines.CoxTimeVaryingFitter. Proportional-hazards assumptions are tested via scaled Schoenfeld residuals; violating covariates are stratified rather than regressed. Tenant is included as a frailty term to absorb unobserved heterogeneity.

### 4.4. In-Platform Scoring

Fitted coefficients are exported as a lookup table and evaluated natively - a Splunk lookup joined in a scheduled search, an ASIM watchlist joined in KQL - so hazard scores are computable without egressing detection telemetry from the tenant boundary. This is the key deployability constraint in multi-tenant environments with data-residency obligations.

### 4.5. Scheduling Service

The knapsack solver runs nightly, emitting a ranked revalidation backlog into the existing engineering ticket queue with hazard score, attributed cause, and estimated effort. Analysts see routine tickets; no new tooling is imposed.




**Fig 2: Mulity Tenant Decay Scoring and Governance Architecture**

The architecture inverts the usual multi-tenant analytics direction. Rather than centralising tenant telemetry for scoring, only the fitted coefficient vector is distributed outward, and scoring executes inside each tenant. Cross-tenant learning is preserved because coefficients are estimated on de-identified aggregate panels, which carry no per-alert content.

The Schoenfeld feedback edge is not decorative. In practice, ingestion-volume z-shift consistently violates

proportional hazards during tenant onboarding, and stratifying rather than regressing it prevents a large spurious hazard inflation across every rule in a newly onboarded estate.

The disposition feedback loop from triage closes the measurement circuit. Analyst dispositions are simultaneously the label source and the operational output, which means measurement quality is bounded by triage discipline - the principal validity threat we discuss in Section 5.

## 5. Results and Comparative Analysis

### 5.1. Rule Half-Life by ATT&CK Tactic

**Table 1: Estimated Median Rule Half-Life and Decay Attribution by Tactic (N = 1,847 Rules, 24 Months)**

| ATT&CK Tactic             | Rules (n) | Median half-life (days) | 95% CI  | Schema-attributable hazard (%) | Adversary-attributable (%) |
|---------------------------|-----------|-------------------------|---------|--------------------------------|----------------------------|
| Defence Evasion           | 214       | 94                      | 81–108  | 22.4                           | 77.6                       |
| Command & Control         | 187       | 118                     | 103–136 | 31.9                           | 68.1                       |
| Execution                 | 241       | 163                     | 148–181 | 44.6                           | 55.4                       |
| Credential Access         | 268       | 207                     | 189–228 | 63.8                           | 36.2                       |
| Lateral Movement          | 196       | 221                     | 198–247 | 52.1                           | 47.9                       |
| Persistence               | 259       | 244                     | 224–268 | 61.2                           | 38.8                       |
| Discovery                 | 233       | 289                     | 261–319 | 58.7                           | 41.3                       |
| Collection / Exfiltration | 249       | 312                     | 284–344 | 55.3                           | 44.7                       |
| All rules                 | 1,847     | 211                     | 203–220 | 51.6                           | 48.4                       |

Half-life varies by a factor of 3.3 across tactics, refuting the uniform-decay assumption embedded in calendar review. Defence-evasion content decays fastest and is overwhelmingly adversary-driven, consistent with evasion being the tactic under most direct optimisation pressure. Credential-access and persistence detections show the inverse

profile: they decay more slowly, but 61–64% of their hazard is schema-attributable, meaning most of their failure is self-inflicted parser and normalisation drift rather than adversary innovation. The aggregate near-even split (51.6% schema) is misleading if applied uniformly - it is the per-tactic decomposition, not the mean, that carries operational value.




Fig 3: Rule Half-Life and Schma Attributable Hazard by Attace and Tactic

## 5.2. Governance Model Comparison

Table 2: Fixed-Interval Review Baseline Versus Decay-Aware Scheduling (12-Month Prospective Evaluation, Equal Budget)

| Metric                                      | Baseline (90-day fixed) | Decay-aware (hazard-ranked) | $\Delta$ | p-value |
|---------------------------------------------|-------------------------|-----------------------------|----------|---------|
| Retained detection efficacy (%)             | 68.1                    | 87.4                        | +19.3 pp | < 0.001 |
| Detection precision (TP / total alerts)     | 0.412                   | 0.594                       | +0.182   | < 0.001 |
| False-positive volume (monthly, indexed)    | 100.0                   | 69.4                        | -30.6%   | < 0.001 |
| Analyst-hours per true positive             | 3.42                    | 2.25                        | -34.2%   | < 0.001 |
| Silent-failure rules undetected $\geq 90$ d | 147                     | 31                          | -78.9%   | < 0.001 |
| Mean time to revalidation (days)            | 88.6                    | 41.3                        | -53.4%   | < 0.001 |
| Engineer-hours consumed (budget)            | 1,440                   | 1,440                       | 0.0      | -       |
| Rules revalidated (n)                       | 1,847                   | 968                         | -47.6%   | -       |

The decay-aware arm achieves superior outcomes on every efficacy metric while revalidating fewer than half as many rules at identical engineer-hours. This is the central finding: calendar review is not merely suboptimal but actively wasteful, spending roughly half its budget on low-hazard content that had not degraded. The 78.9% reduction in long-lived silent failures is the most operationally consequential

result, since silent failures are undetectable by false-positive-driven governance and represent uncompensated coverage loss. Paired *t*-tests across the eleven tenants confirm significance at  $p < 0.001$  for all efficacy endpoints; the effect direction was consistent in every tenant, with effect magnitude greatest in tenants undergoing active log-source migration.




Fig 4: Governance Model Comparison at Equal Revalidation Budget

### 5.3. Co5.2. Model Covariate Effects

**Table 3: Cox Proportional-Hazards Coefficients (Time-Varying, Tenant Frailty; Concordance = 0.783)**

| Covariate                                   | $\beta$ | Hazard ratio | 95% CI (HR) | p-value |
|---------------------------------------------|---------|--------------|-------------|---------|
| CIM/ASIM conformance failure rate (per 10%) | 0.641   | 1.898        | 1.71–2.11   | < 0.001 |
| Schema-change event (30-day lag)            | 0.522   | 1.685        | 1.49–1.91   | < 0.001 |
| Index-time parsing error rate (per 5%)      | 0.398   | 1.489        | 1.32–1.68   | < 0.001 |
| Ingestion-volume z-shift (per $1\sigma$ )   | 0.214   | 1.239        | 1.11–1.38   | < 0.001 |
| Tactic = Defence Evasion                    | 0.487   | 1.627        | 1.38–1.92   | < 0.001 |
| Data source = EDR telemetry                 | 0.176   | 1.192        | 1.04–1.37   | 0.012   |
| Data source = Cloud audit (multi-cloud)     | 0.309   | 1.362        | 1.18–1.57   | < 0.001 |
| Cumulative tuning interventions (per event) | −0.288  | 0.750        | 0.68–0.83   | < 0.001 |
| Rule authored under detection-as-code       | −0.412  | 0.662        | 0.58–0.76   | < 0.001 |
| Tenant vertical = Financial services        | −0.093  | 0.911        | 0.79–1.05   | 0.191   |

Conformance failure is the single strongest hazard driver, nearly doubling instantaneous failure risk per ten-percentage-point increase - a metric most SOCs treat as platform hygiene rather than detection risk. Schema-change events remain significant at 30-day lag, confirming that degradation follows drift with a measurable delay long enough to permit pre-emptive scheduling. Two protective effects are notable: each tuning intervention reduces hazard by 25%, and detection-as-

code authorship confers a durable 33.8% hazard reduction, evidence that engineering rigour at authoring time yields compounding longevity rather than a one-off quality gain. Tenant vertical is non-significant ( $p = 0.191$ ), indicating decay is a property of content and telemetry rather than of industry threat profile - the finding that justifies cross-tenant coefficient sharing in Figure 2.




**Fig 5: Cox Hazard Ratios for Detection Decay Covariates**

## 6. Conclusion

This work establishes detection decay as a measurable, predictable, and governable property of deployed SIEM content, contributing the first survival-analytic estimator of rule half-life and the first decomposition of decay into schema-induced and adversary-induced hazard. Across 1,847 rules and eleven tenants, median half-life was 211 days with a 3.3-fold spread by ATT&CK tactic, and slightly over half of aggregate hazard proved self-inflicted rather than adversarial - a reallocation signal invisible to coverage-matrix governance. Converting hazard into a bounded scheduling policy retained 87.4% of detection efficacy against 68.1% for calendar review, cut analyst-hours per true positive by 34.2%, and reduced long-lived silent failures by 78.9% at unchanged

engineering cost. Practically, the method deploys on telemetry that mature SOCs already emit and scores rules inside the tenant boundary, satisfying data-residency constraints. Immediate next steps are threefold: extend estimation to competing-risks models that treat retirement and evasion as distinct failure modes; incorporate threat-intelligence covariates to anticipate adversary-driven hazard before yield declines; and validate the coefficient-transfer assumption on estates outside the originating MSSP. Longer term, half-life should become a published attribute of shared detection content, allowing consumers of community rule libraries to reason about expected lifetime as routinely as they reason about coverage.

## References

- [1] Bhatt, S., Manadhata, P. K., & Zomlot, L. (2018). The operational role of security information and event management systems. *IEEE Security & Privacy*, 16(5), 35–41.
- [2] Chen, Y., Zhang, Y., & Liu, X. (2020). Concept drift detection in streaming security analytics. *Computers & Security*, 95, 101869.
- [3] Sandeep Kamadi. (2022). Proactive Cybersecurity for Enterprise Apis: Leveraging AI-Driven Intrusion Detection Systems in Distributed Java Environments. *International Journal of Research in Computer Applications and Information Technology (IJRCAIT)*, 5(1), 34-52. [https://iaeme.com/MasterAdmin/Journal\\_uploads/IJRC AIT/VOLUME\\_5\\_ISSUE\\_1/IJRCAIT\\_05\\_01\\_004.pdf](https://iaeme.com/MasterAdmin/Journal_uploads/IJRC AIT/VOLUME_5_ISSUE_1/IJRCAIT_05_01_004.pdf)
- [4] inque, M., Cotroneo, D., & Pecchia, A. (2018). Challenges and directions in security information and event management. *Proceedings of the IEEE International Symposium on Software Reliability Engineering Workshops*, 95–99.
- [5] Feng, C., Wu, S., & Liu, N. (2017). A user-centric machine learning framework for cyber security operations centers. *Proceedings of the IEEE International Conference on Intelligence and Security Informatics*, 173–175.
- [6] Sivaramakrishnan Narayanan (2022). Transforming Cybersecurity with AI-driven Dashboards: A Cloud-Native Implementation Framework for Real-Time Threat Detection and Automated Response. *International Journal of Future Innovative Science and Technology (IJFIST)*, Vol. 5 No. 5 (2022): *International Journal of Future Innovative Science and Technology (IJFIST)*, pp. 9207-9217. <https://doi.org/10.15662/IJFIST.2022.0505004>
- [7] Hassan, W. U., Guo, S., Li, D., Chen, Z., Jee, K., Li, Z., & Bates, A. (2019). NoDoze: Combatting threat alert fatigue with automated provenance triage. *Proceedings of the Network and Distributed System Security Symposium*.
- [8] Lakshmi Kiran Meesala "AI-Driven Cyber Defense: A Hybrid Deep Learning Framework for Real-Time Threat Prediction and Response" *International Journal of Scientific Research in Science, Engineering and Technology (IJSRSET)*, Print ISSN : 2395-1990, Online ISSN : 2394-4099, Volume 10, Issue 2, pp.916-930, March-April-2023. Available at doi : <https://doi.org/10.32628/IJSRSET235845>
- [9] Kottenko, I., & Chechulin, A. (2017). Analytical modeling and assessment of security event correlation. *Journal of Computer Virology and Hacking Techniques*, 13(3), 187–199.
- [10] Kwon, D., Kim, H., Kim, J., Suh, S. C., Kim, I., & Kim, K. J. (2019). A survey of deep learning-based network anomaly detection. *Cluster Computing*, 22, 949–961.
- [11] Legoy, V., Caselli, M., Seifert, C., & Peter, A. (2020). Automated retrieval of ATT&CK tactics and techniques for cyber threat reports. *arXiv preprint arXiv:2004.14322*.
- [12] Nadeem, A., Verwer, S., Moskal, S., & Yang, S. J. (2021). Alert-driven attack graph generation using S-PDFA. *IEEE Transactions on Dependable and Secure Computing*, 19(2), 731–746.
- [13] Sandeep Kamadi, " Risk Exception Management in Multi-Regulatory Environments: A Framework for Financial Services Utilizing Multi-Cloud Technologies" *International Journal of Scientific Research in Computer Science, Engineering and Information Technology (IJSRCSEIT)*, ISSN : 2456-3307, Volume 7, Issue 5, pp.350-361, September-October-2021. Available at doi : <https://doi.org/10.32628/CSEIT217560>
- [14] Oprea, A., Li, Z., Norris, R., & Bowers, K. (2018). MADE: Security analytics for enterprise threat detection. *Proceedings of the Annual Computer Security Applications Conference*, 124–136.
- [15] Meesala, A. (2022). Adaptive Spread Anomaly Intelligence Framework (ASAI): A cloud-native AI framework for real-time bid-ask spread anomaly detection and cross-venue liquidity risk intelligence. *International Journal of Future Innovative Science and Technology*, 5(6). <https://doi.org/10.15662/IJFIST.2022.0506007>
- [16] Rahman, M. R., Mahdavi-Hezaveh, R., & Williams, L. (2020). A literature review on mining cyberthreat intelligence from unstructured texts. *Proceedings of the IEEE International Conference on Data Mining Workshops*, 516–525.
- [17] Sivaramakrishnan Narayanan (2023). Operationalizing Artificial Intelligence Security in the Cloud: A Practical Integration framework for Enterprise Risk Management. *International Journal of Future Innovative Science and Technology (IJFIST)*, Vol. 6 No. 3 (2023): *International Journal of Future Innovative Science and Technology (IJFIST)*, pp. 10611-10619. <https://doi.org/10.15662/IJFIST.2023.0603002>
- [18] Strom, B. E., Applebaum, A., Miller, D. P., Nickels, K. C., Pennington, A. G., & Thomas, C. B. (2018). *MITRE ATT&CK: Design and philosophy*. MITRE Corporation.
- [19] Vielberth, M., Böhm, F., Fichtinger, I., & Pernul, G. (2020). Security operations center: A systematic study and open challenges. *IEEE Access*, 8, 227756–227779.
- [20] Sandeep Kamadi, " Adaptive Federated Data Science & MLOps Architecture: A Comprehensive Framework for Distributed Machine Learning Systems" *International Journal of Scientific Research in Computer Science, Engineering and Information Technology(IJSRCSEIT)*, ISSN : 2456-3307, Volume 8, Issue 6, pp.745-755, November-December-2022. Available at doi : <https://doi.org/10.32628/CSEIT22555>
- [21] Wagner, T. D., Mahbub, K., Palomar, E., & Abdallah, A. E. (2019). Cyber threat intelligence sharing: Survey and research directions. *Computers & Security*, 87, 101589.
- [22] Lakshmi Kiran Meesala, " Modern Security Information and Event Management: Architecture, Analytics Pipelines, and Empirical Evaluation of SOC-Scale Threat Detection" *International Journal of Scientific Research in Computer Science, Engineering and Information Technology(IJSRCSEIT)*, ISSN : 2456-3307, Volume 9, Issue 4, pp.995-1012, July-August-2023. Available at doi : <https://doi.org/10.32628/CSEIT23564538>

- [22] Zimmermann, M., Gerhardt, A., & Kaiser, D. (2021). Reliability-centred maintenance scheduling using proportional hazards modelling. *Reliability Engineering & System Safety*, 209, 107470.
- [23] Arun Meesala , " A Distributed Kafka-Centric Framework for High-Throughput Mid-Price Computation and Intelligent Time-Series Persistence in Financial Clouds" International Journal of Scientific Research in Computer Science, Engineering and Information Technology(IJSRCSEIT), ISSN : 2456-3307, Volume 9, Issue 2, pp.998-1006, March-April-2023. Available at doi : <https://doi.org/10.32628/CSEIT2342442>.