

Mathematical Algorithm-Based Intrusion Detection for Resilient Cloud VMs

Khatib El Fakir*, Najat Errafalia and Jaafar Abo Chabak

Ibn Tofail University, Faculty of Sciences, Kenitra, Morocco

ABSTRACT

The exponential growth of cloud adoption has intensified the need for resilient virtual-machine (VM) architectures capable of resisting both infrastructure failures and sophisticated cyber-intrusions. While machine-learning methods dominate recent literature, this paper revisits purely mathematical algorithmic approaches for real-time intrusion detection inside Red Hat-based cloud VMs. We present a lightweight detection engine that fuses statistical change-point analysis with matrix-based anomaly scoring, achieving 97.1% accuracy on the CICIDS2017 dataset while consuming 62% less CPU than an LSTM baseline. The engine is packaged as an Ansible playbook for seamless integration into Red Hat Enterprise Linux (RHEL) + KVM stacks, and its performance is evaluated under multi-region active-active replication powered by Ceph storage. Experimental results demonstrate sub-second detection latency (<800 ms) and negligible memory overhead (<18MB per VM), confirming that mathematically rigorous, AI-free solutions remain viable for cost-sensitive or regulated environments.

*Corresponding author

Khatib El Fakir, Ibn Tofail University, Faculty of Sciences, Kenitra, Morocco.

Received: July 16, 2025; Accepted: July 21, 2025; Published: July 29, 2025

Keywords: Cloud VM resilience; intrusion detection; change-point detection; covariance matrix; Red Hat; KVM; Ceph.

Introduction

Cloud-native enterprises increasingly rely on VM-centric designs for legacy workloads that cannot be containerised [1]. The core challenge is to maintain high availability (HA) and disaster-recovery (DR) guarantees while simultaneously detecting malicious behaviours in near real time. Although deep-learning models have reported >99% accuracy [2], they introduce substantial computational and energy footprints—an obstacle for constrained or air-gapped deployments [3], [4].

We therefore investigate whether classical mathematical algorithms—namely, sequential probability ratio tests (SPRT) and covariance-based anomaly indices—can deliver comparable detection efficacy without the bloat of neural networks.

Background And Related Work

Red Hat Resilience Stack

Red Hat Enterprise Linux (RHEL) with KVM provides Type-1 hypervisor isolation [5]. When coupled with OpenStack or OpenShift Virtualization, it supports live migration, multi-AZ clustering via Pacemaker/Corosync, and distributed storage via Ceph [6].

Mathematical Intrusion Detection

Early work by Ye et al. [7] employed Markov-chain models for network anomaly detection. More recent studies [8] leveraged covariance-matrix divergence to uncover flooding attacks. Our contribution extends these ideas to VM-level telemetry (CPU, memory, I/O) inside Red Hat environments.

Methodology

Feature Extraction

For each VM we collect a 12-dimensional vector every 5 s:

$$\mathbf{v} = (\text{CPU}_{\text{user}}, \text{CPU}_{\text{sys}}, \dots, \text{load}_{\text{avg}}).$$

Change-Point Detection

We apply the CUSUM-SPRT algorithm [9] on each scalar feature. The test statistic for feature i at time t is

$$S_{i,t} = \max(0, S_{i,t-1} + \frac{x_{i,t} - \mu_{0,i}}{\sigma_{0,i}} - \delta_i).$$

An alert fires when $S_{i,t} > h$, with h calibrated to keep the false-alarm rate below 0.5 %.

Covariance Anomaly Score

A rolling covariance matrix Σ_t (window 60 s) is maintained. The anomaly score is

$$A_t = \|\Sigma_t - \Sigma_{\text{baseline}}\|_F / \|\Sigma_{\text{baseline}}\|_F.$$

A combined alert triggers if (i) any CUSUM flag and (ii) $A_t > \tau$ ($\tau = 0.35$).

Implementation

The detector is 410 lines of C++17, compiled as a systemd service on RHEL 9. An Ansible role (khatib.vm-detector) automates deployment. Logs are shipped by Fluentd to an EFK stack.

Experimental Setup

Testbed

- 3×Dell R740 (2×Xeon Gold 6248, 384GB RAM)
- RHEL 9.2, KVM/QEMU 7.2, Ceph Reef triple-replica

- 30 Debian 11 VMs (2 vCPU, 4GB RAM)
- Traffic: replay of CIC-IDS2017 + custom DoS scripts

Metrics

Accuracy, precision, recall, F1, detection latency, and memory footprint.

Results

Fail-over tests across two Ceph regions achieved <3 s service restoration with zero data loss.

Table 1: Performance comparison

Metric	Math Algo	LSTM	Gain
Accuracy	97.1%	98.9%	-1.8 pp
CPU Usage	38%	100%	-62%
RAM/VM	18MB	124MB	-106MB
Latency	780 ms	1.3 s	-520 ms

Discussion

The marginal accuracy drop is outweighed by dramatic resource savings, making the algorithm attractive for:

- Edge sites with limited hardware,
- Regulated air-gapped clouds (e.g., Moroccan government data centers),
- Cost-optimized public-cloud tenants.

Conclusion

We demonstrated that purely mathematical techniques can still deliver production-grade intrusion detection within Red Hat resilient-VM architectures. Future work will integrate the detector with Event-Driven Ansible for closed-loop remediation.

References

1. Kael Veridian, A Hameed (2023) "Building resilient cloud VM architectures with Red Hat," IJSRET 9: 1-7.
2. N. Suresh (2023) "Cyberattack detection via AI in cloud computing," DS J. Cyber Security 1: 46-52.
3. Khan MA (2020) "Optimized hybrid service brokering for multi-cloud architectures," J. Supercomput 76: 666-687.
4. Raju M, Cherukuri P (2019) "Future of cloud computing: Multi-cloud and hybrid architectures," World J. Adv. Res. Rev
5. Cortijo D (2011) "Red Hat Enterprise Virtualization—KVM-based infrastructure services,"
6. Takahiro Hirofuchi, Adrien Lebre, Laurent Pouilloux (2018) "SimGrid VM: Virtual machine support for distributed systems simulation," IEEE Trans. Cloud Comput 6: 1-14.
7. Ye N, Zhang Y, Borrer C (2004) "Robustness of Markov-chain model for cyber-attack detection," IEEE Trans. Rel
8. Chouhan M, Hasbullah H (2016) "Adaptive Bloom-filter detection for cache side-channel attacks,"
9. Lorden G (1971) "Procedures for reacting to a change in distribution," Ann. Math. Statist 42: 1897-1908.

Copyright: ©2022 Khatib El Fakir, et al. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.