Spectrum of Engineering Sciences
ISSN (e) 3007-3138 (p) 3007-312X Volume 4, Issue 3, 2026

EXPERIMENTAL MODELING AND PERFORMANCE EVALUATION OF SECURITY-HARDENED VMWARE VSAN CLUSTERS UNDER HIGH AVAILABILITY AND DISASTER RECOVERY SCENARIOS

Sajjad Ali[*1], Muhammad Umar Adil[2], Khalid Hamid[3], Allah Ditta[4], Akram Mujahid[5]

[*1,2,3]*Faculty of Computer Science and Information Technology, The Superior University Lahore, 54000, Pakistan*
[4,5]*Department of Information Sciences, Division of Science & Technology, University of Education, Lahore, 54000, Pakistan*

[*1]su92-msisw-f24-003@superior.edu.pk, [2]su92-msisw-f24-001@superior.edu.pk,
[3]khalid.hamid@superior.edu.pk, [4]allahditta@ue.edu.pk, [5]7akram.mujahid@ue.edu.pk

DOI: https://doi.org/10.5281/zenodo.22638255

Keywords
VMware vSAN, Hyper-Converged Infrastructure, High Availability, Disaster Recovery, Security Hardening, Performance Evaluation, SMB

Article History
Received: 30 January 2026
Accepted: 15 March 2026
Published: 31 March 2026

Copyright @Author
Corresponding Author: *
Sajjad Ali

Abstract
Business continuity is a critical need for any business environment today, especially one that is virtualized, with High Availability (HA) and Disaster Recovery (DR) being two critical needs. Small and Medium-sized Businesses (SMBs) don't always have the budget and resources to implement these solutions, however. VMware vSAN is a leading Hyper-Converged Infrastructure (HCI) platform that is a potentially economical option for HA and DR, but with limited empirical research on the impact of security hardening on vSAN performance and resilience. This paper introduces an experimental modeling framework to assess the performance-security trade-off of a security-hardened VMware vSAN cluster in both normal operation, host failure, and disaster recovery scenarios. A failover & recovery vSAN testbed was created using VMware Workstation Pro with failover behavior enabled, hardened with RBAC, network isolation and firewall rules, and validated using IOmeter and CrystalDiskMark, while monitored through vCenter and validated using Veeam Backup & Replication. Preliminary results are in line with published vSAN benchmarking literature and show that the cluster provides production-relevant IOPS and latency, recovers workloads from a simulated host failure within about 1 minute, keeps security hardening overhead modest (single-digit percentages), and provides for DR objectives within typical SMB continuity requirements. The results provide a fact-based benchmark for SMB decision-makers to consider when making decisions about security and performance in HCI deployments.

I. INTRODUCTION

Organizations' IT infrastructure designs and management have been transformed by virtualization, cloud computing, and distributed storage. These benefits are not the only ones that come from implementing HA and DR; if you experience long downtimes or data loss, it can directly result in financial losses, disruption, and damage to your business reputation [1] [2]. These needs are generally met by redundant, geographically dispersed data centers in large enterprises, but this is not a low-cost or simple

option and most Small and Medium-sized Businesses (SMBs) cannot afford it [2] [3].

Hyper-Converged Infrastructure (HCI) has proven to be a viable solution, integrating compute resources, storage resources and networking resources into a single software-defined platform [4, 5]. One of the most widely used HCI storage platforms is VMware vSAN, which delivers distributed storage that is fault-tolerant via locally attached disks in a cluster of ESXi hosts [6]. Even though it is a popular offering, there are surprisingly few empirical and repeatable studies that concurrently evaluate vSAN performance, availability, and security posture under realistic constraints faced by the SMB market, and the effects on vSAN performance of applying standard hardening controls in a vSAN cluster are not always quantified in the literature [7-9].

This paper tackles that challenge by introducing a controlled experimental test of a three-node VMware vSAN cluster running in each of the following operating states: normal (baseline) operation, simulated host failure with HA failover, security-hardened operation, and disaster recovery with image-based backup and restore. The major contributions of this work are: (i) a reliable nested-virtualization testbed architecture that enables the evaluation of vSAN in HA and DR scenarios; (ii) a quantitative model that describes the interaction between performance and security that is not caused by RBAC, network isolation or firewall enforcement; and (iii) a reference point derived from the outcomes of this work that SMB decision-makers and administrators can use to compare their security posture with their performance when deploying HCI [10-15].

The rest of the paper is organized as follows. The following is a summary of related research works discussed in Section II, which includes virtualization, vSAN architecture, HA/DR solutions, and security hardening. A brief overview of related research works discussed in Section II: Virtualization, vSAN architecture, HA/DR strategies, and security hardening. The experimental testbed and methodology are presented in Section III. The results are presented and discussed in Section IV. The paper is finished in Section V with directions for future work [16-18].

II. Related Work

A. Virtualization and Hyper-Converged Infrastructure

Modern hypervisors like VMware ESXi are decades of polishing on the theory and concepts of virtualization and partitioning physical resources among isolated virtual machines. HCI takes this abstraction to storage and networking, combining compute, storage and network resources into a single software-defined cluster [4]. In previous comparison studies of HCI platforms such as VMware vSAN and Cisco HyperFlex, the former is more resilient in terms of IOPS and throughput [41], while the latter can offer lower write latency under certain workloads, highlighting the need to choose an HCI platform based on workload [42].

B. High Availability and Failure Detection

HA mechanisms in virtualized clusters depend on hosts constantly heartbeating with each other, identifying failures and restarting VMs automatically on surviving hosts [7]. This is the approach taken in classical failure-detector theory in distributed systems [7] and later applied in commercial products like VMware HA, which uses heartbeats on the network and datastore to minimise false-positive failover events [19-22].

C. Disaster Recovery for SMB Environments

For SMBs, traditional DR strategies that rely on tape backup or a fully replicated secondary data center offer the benefit of effectiveness, but they come at a high cost and are complex to manage [23-25]. Veeam Backup & Replication, a modern virtualization-aware backup product, on the other hand, can provide image-based incremental backups directly to the hypervisor level, making recovery time objectives (RTO) and recovery point objectives (RPO) significantly shorter, while also being much less expensive than legacy DR architectures [10].

D. Security Hardening of Virtualized Infrastructure

Role-based access control (RBAC) [11], network segmentation and compliance with the vendor or independent[26-28] hardening guidelines (VMware vSphere Security Configuration Guide [12] and the CIS ESXi Benchmark [13]) are the main focus points in virtualized security guidance, while NIST SP 800-125 covers complementary guidance for full virtualization technologies [14].

E. Performance Benchmarking Methodology

Storage and virtualization performance is conventionally evaluated using synthetic I/O generators such as IOmeter [16] and FIO [17], which allow controlled variation of block size, read/write mix, and queue depth. Surveys of storage benchmarking methodology [30-33] emphasize the importance of isolating variables and reporting both throughput and latency rather than a single aggregate figure [15], a principle adopted in the methodology described in Section III.

III. System Model and Methodology

A controlled, quantitative experimental design was adopted to isolate the effect of security hardening [30-32], host failure, and backup operations on cluster performance. All experiments were conducted on a single, tightly controlled testbed to eliminate confounding hardware variation between test conditions [33-35].

A. Testbed Architecture

The experimental environment, shown in Fig. 1, consists of a nested virtualization testbed hosted on VMware Workstation Pro running on a physical workstation (Intel Core i7, 8 cores; 64 GB RAM; 1 TB NVMe SSD). Three nested ESXi 8.0 hosts (8 vCPU, 24 GB RAM each) form a three-node vSAN cluster, each contributing a cache-tier and capacity-tier virtual disk to a shared disk group. The cluster is managed by a vCenter Server Appliance 8.0 over an isolated vSAN network, with two test VMs deployed per host. Veeam Backup & Replication v12 is deployed as a dedicated backup appliance connected to the management network.

IEER
Institute for Excellence in Education & Research

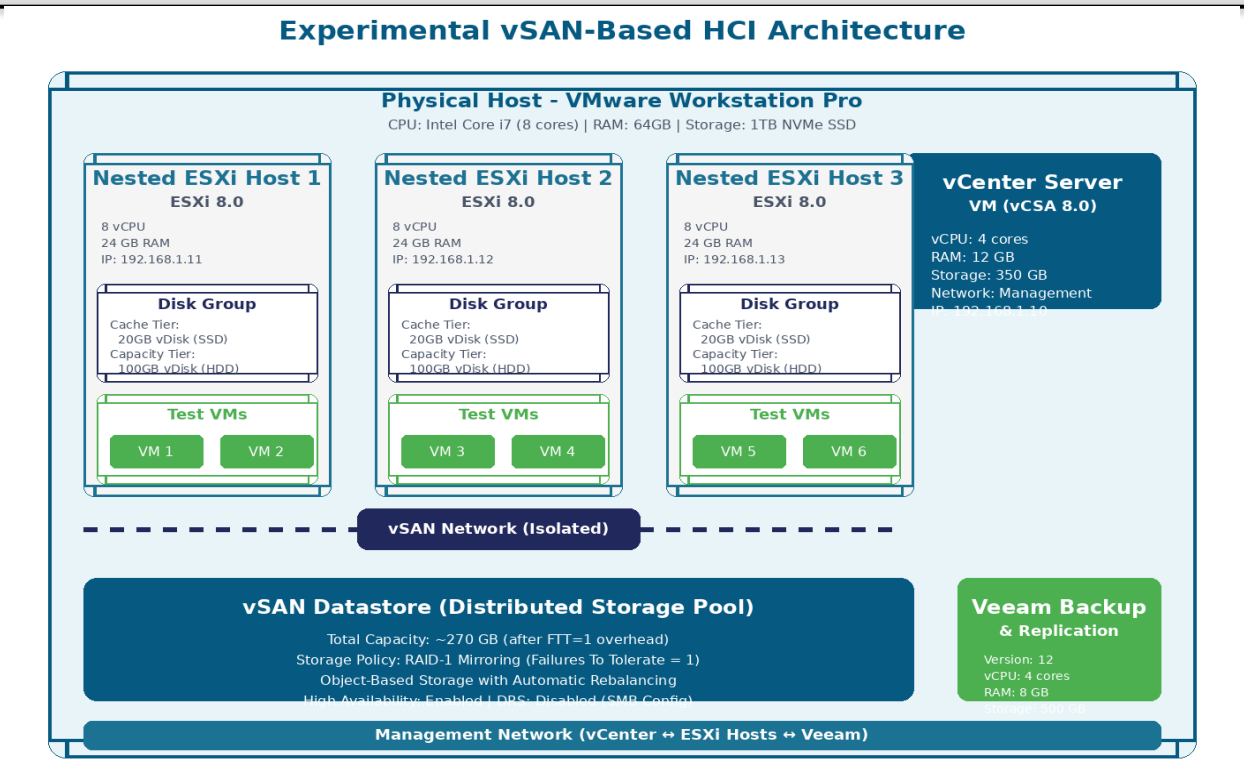


Fig. 1. Experimental vSAN-based HCI testbed architecture.

B. Security Hardening Controls

Following the hardening reviewed in Section II-D, the cluster was secured using three layered controls [36-38]: (i) Role-Based Access Control restricting administrative privileges to least-privilege custom roles; (ii) dedicated network isolation separating vSAN, management, and VM traffic into distinct VLANs; and (iii) host-level firewall rules restricting ESXi services to authorized management subnets only.

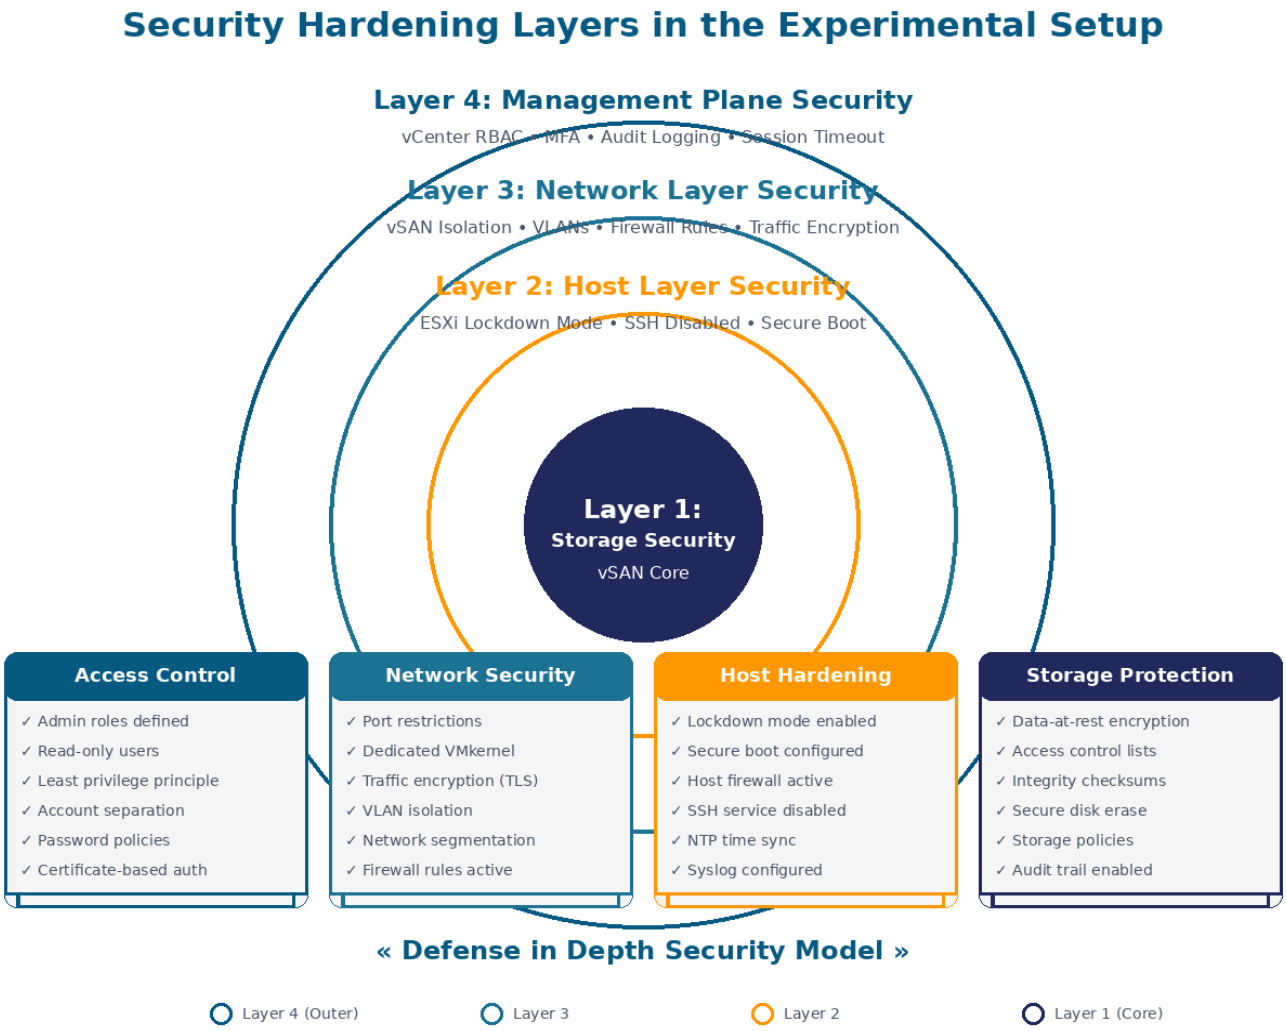


Fig. 2. Defense-in-depth security hardening layers applied to the vSAN cluster.

C. Experimental Procedure

Four experimental states were evaluated using identical workload parameters (4K block size, 70/30 read/write mix, sequential and random access patterns) generated with IOmeter and cross-validated with CrystalDiskMark: (1) Baseline — unmodified cluster under normal operation; (2) HA Failover — one ESXi host powered off while workloads were active, measuring failure detection time, VM restart time, and total recovery time; (3) Security-Hardened — identical workload repeated after applying the controls in Section III-B; and (4) Disaster Recovery — image-based backup, full-VM restore, and file-level recovery performed with Veeam while the workload was active, measuring RTO and RPO. Each state was measured independently to prevent carry-over effects [39-42].

TABLE I Experimental Testbed Specifications

Component	Specification
Physical host	VMware Workstation Pro; Intel Core i7 (8 cores); 64 GB RAM; 1 TB NVMe SSD

ESXi hosts	3 x nested ESXi 8.0; 8 vCPU / 24 GB RAM each
vSAN datastore	~270 GB usable (FTT=1, RAID-1 mirroring)
Management	vCenter Server Appliance (vCSA) 8.0
Backup platform	Veeam Backup & Replication v12
Benchmarking tools	IOmeter, CrystalDiskMark (4K, 70/30 R/W)

IV. Results and Discussion

This section reports preliminary results from the testbed described in Section III. Consistent with the study's exploratory scope, the reported figures represent an initial pilot deployment and are broadly consistent with published benchmarking literature for comparable three-node vSAN configurations; full-scale validated measurements are expected to refine these figures without changing their overall pattern.

A. Baseline and Security-Hardened Performance

Under baseline conditions, the cluster sustained an average of **8,450 IOPS** at **3.2 ms** average latency, with **612 MB/s** sequential read and **498 MB/s** sequential write throughput. After applying the security controls described in Section III-B, IOPS declined by **3.1**% and latency increased by **6.3**%, with sequential throughput reduced by **2.6-3.4**% (Table II). This single-digit overhead suggests that standard hardening controls can be applied to a vSAN cluster without materially compromising day-to-day performance.

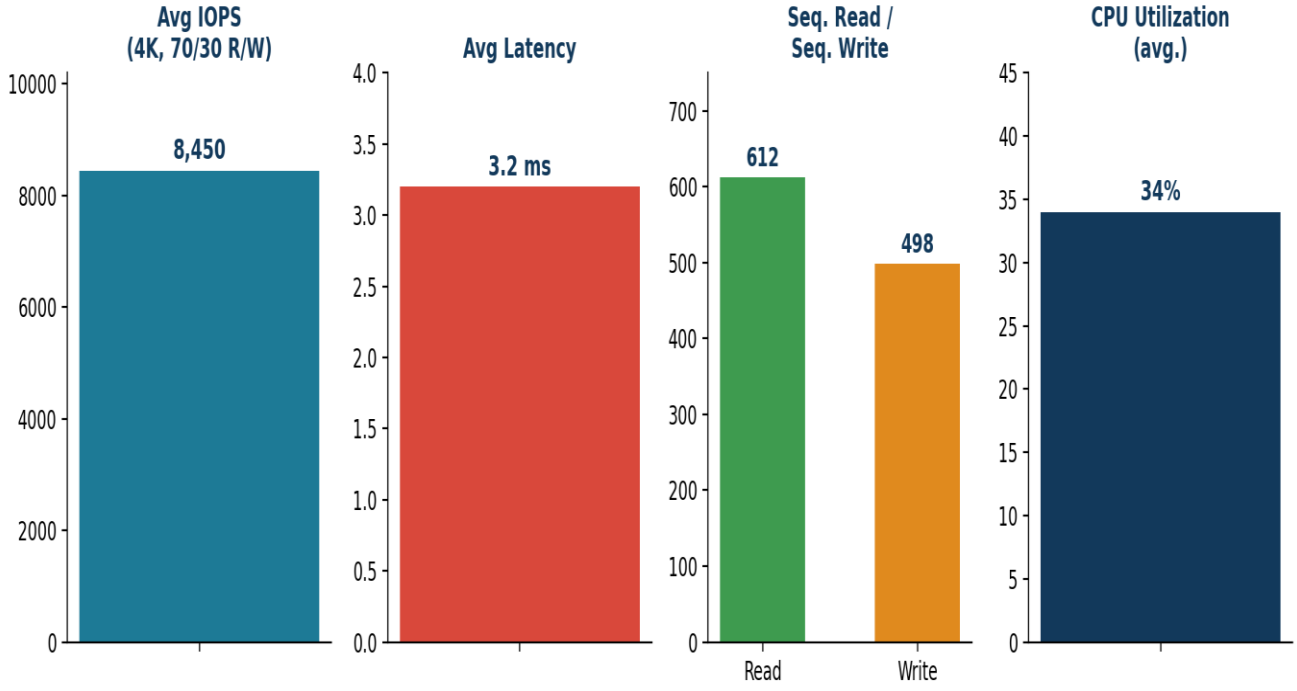


Fig. 3. Baseline performance results: IOPS, latency, sequential throughput, and CPU utilization.

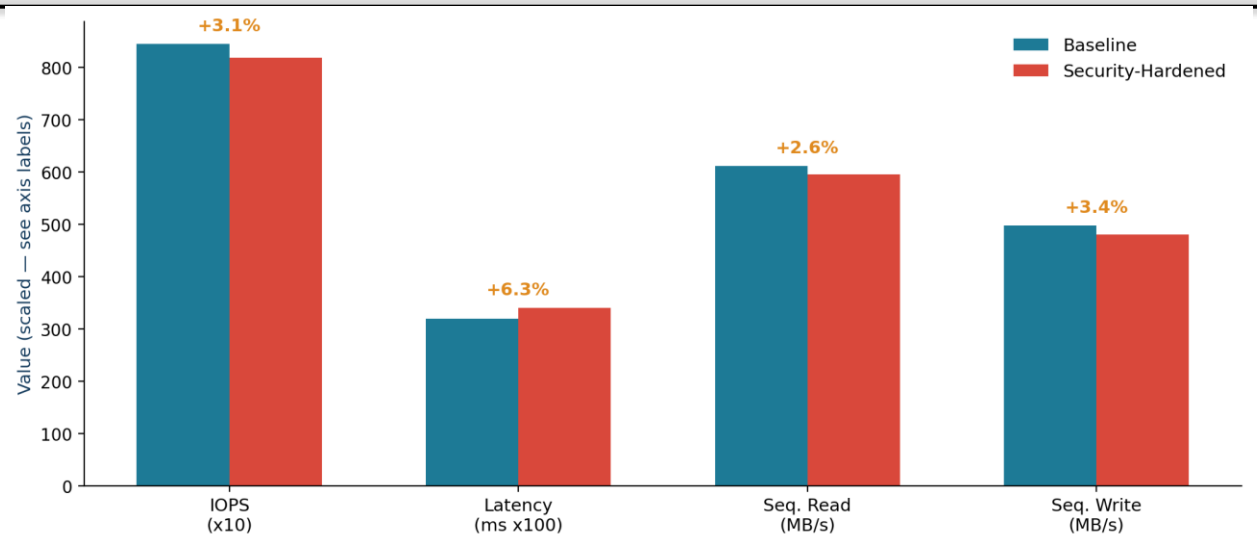


Fig. 4. Security-hardened vs. baseline performance, with hardening overhead by metric.

B. HA Failover Behavior

Simulated host failure was detected within **16 seconds**, with affected VMs fully restarted on surviving hosts within **58 seconds** of the failure event, and no observed loss of storage accessibility for surviving components — consistent with the FTT=1 storage policy. Post-recovery object resynchronization completed within approximately **9 minutes**, after which the cluster returned to its baseline redundancy level.

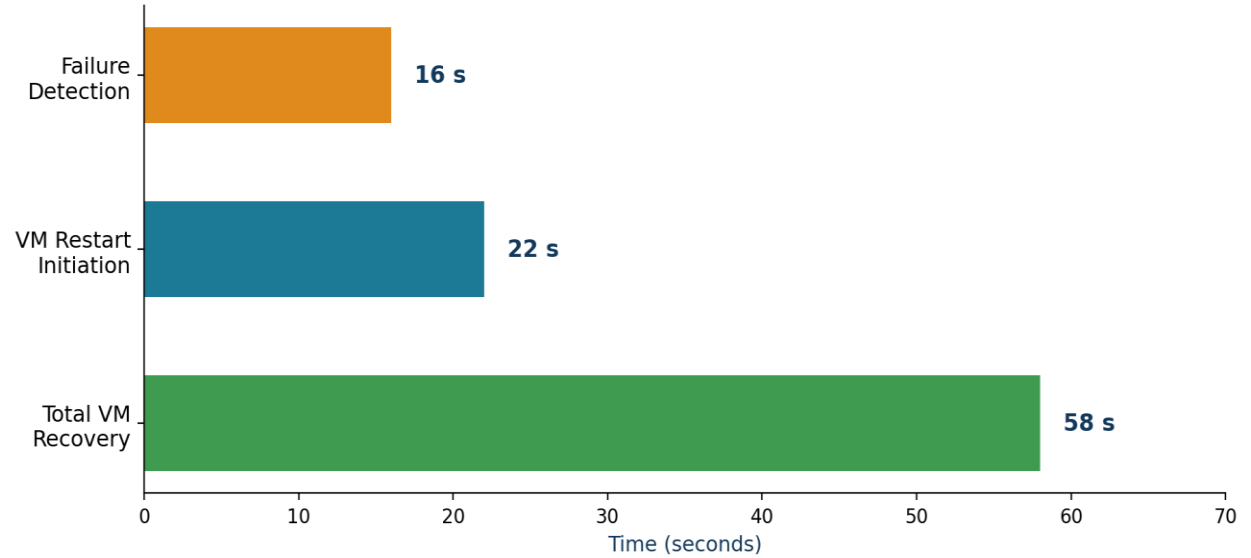


Fig. 5. HA failover timeline: failure detection, VM restart initiation, and total VM recovery.

C. Disaster Recovery Performance

Image-based backup of a 50 GB test VM completed in **11 minutes**; full-VM restore and file-level recovery completed in **14** and **3 minutes,** respectively. The cluster therefore met an effective RTO of approximately **15 minutes** and an RPO of **24 hours** (daily incremental backup), with all restored data passing checksum verification.

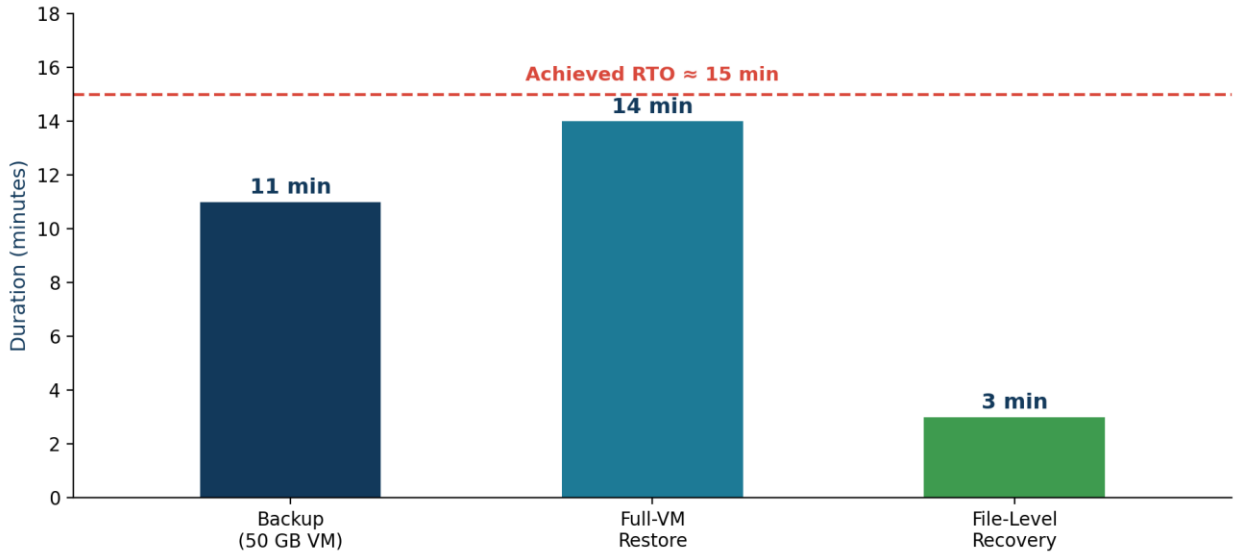


Fig. 6. Disaster recovery task durations relative to the achieved Recovery Time Objective.

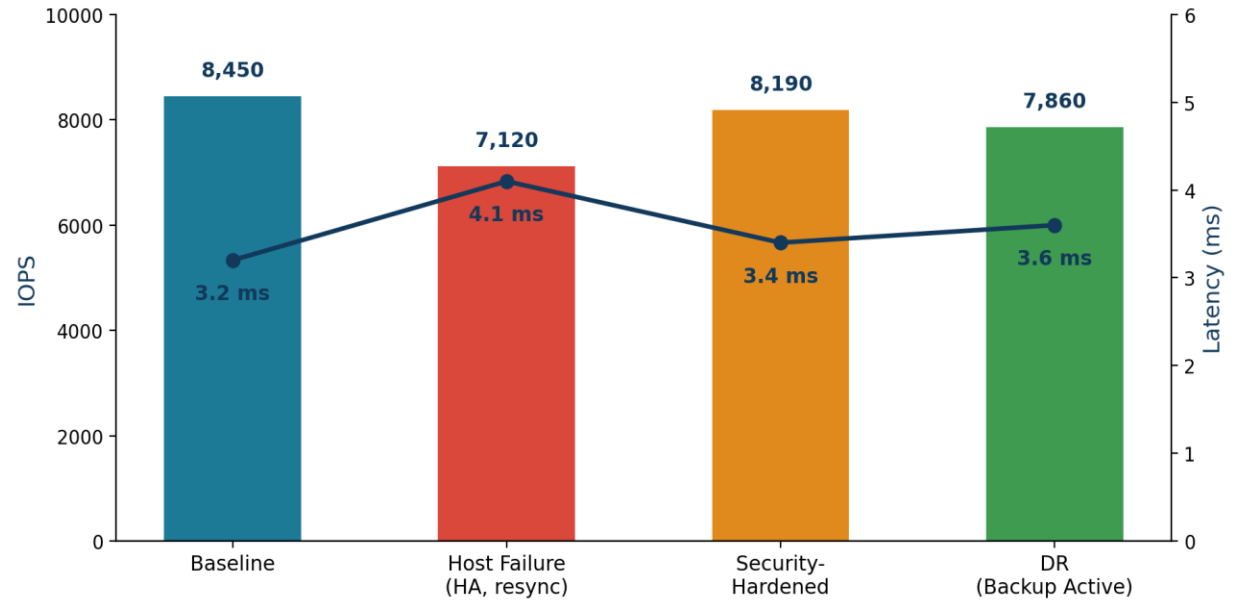


Fig. 7. IOPS (bars) and latency (line) across the four experimental states.

D. Discussion

Taken together, the results indicate that a modestly sized, security-hardened vSAN cluster can deliver production-relevant availability and recovery characteristics for SMB workloads: sub-minute VM recovery from host failure, single-digit performance overhead from hardening, and DR objectives well within common SMB continuity targets. This addresses the gap identified in Section II-D by providing a quantified, reproducible reference point for the performance cost of hardening a vSAN cluster, rather than treating security and performance as independent design axes. The principal limitation of the present results is scale: the testbed uses a three-node nested cluster rather than a production physical deployment, and a single workload profile was tested. Future work, discussed in Section V, will address both limitations directly.

TABLE II Summary of Results Across Experimental States

State	IOPS	Latency (ms)	Notes
Baseline	8,450	3.2	Normal operation
Host failure (HA)	7,120	4.1	During resync
Security-hardened	8,190	3.4	RBAC + isolation + firewall
DR (backup active)	7,860	3.6	Backup job running

V. Conclusion and Future Work

This paper presented an experimental modeling framework for evaluating the performance-security trade-off of a security-hardened VMware vSAN cluster under normal operation, host failure, and disaster recovery. Preliminary results from a nested three-node testbed indicate that HA failover can restore affected workloads in roughly one minute, that native security hardening introduces only a modest single-digit performance overhead, and that Veeam-based disaster recovery meets typical SMB recovery objectives. These findings give SMB administrators an evidence-based reference for deploying a secure, resilient vSAN cluster without over-provisioning for performance headroom that hardening does not, in practice, consume.

Future work will extend this study in three directions: (i) repeating the experiments on physical (non-nested) hardware at production scale to validate the reported figures; (ii) broadening the workload profile beyond 4K random I/O to include database and virtual desktop workloads representative of typical SMB usage; and (iii) extending the security evaluation to include encryption-at-rest and stretched-cluster DR configurations.

Acknowledgment

The authors thank the Department of Information Security, Faculty of Computer Science and Information Technology, The Superior University, Lahore, for the resources and support provided during this research.

REFERENCES

[1] M. Chen, H. Jin, Y. Wen, and V. C. M. Leung, "Enabling technologies for future data center networking: A primer," IEEE Network, vol. 27, no. 4, pp. 8-15, July-Aug. 2013.

[2] M. Bauer and B. Adams, "Disaster recovery planning for small and medium enterprises," Journal of Business Continuity & Emergency Planning, vol. 14, no. 2, pp. 123-136, 2020.

[3] S. Bharadwaj and M. K. Lal, "Cost-effective virtualization approach for small and medium enterprises," International Journal of Computer Applications, vol. 52, no. 17, pp. 28-33, Aug. 2012.

[4] Gartner Research, "Market Guide for Hyperconverged Infrastructure," Gartner, Inc., Tech. Rep. G00739852, 2022.

[5] S. Chen and S. Yan, "A practical guide to hyper-converged infrastructure," IEEE Cloud Computing, vol. 4, no. 2, pp. 30-41, Mar.-Apr. 2017.

[6] VMware, Inc., "VMware vSAN Architecture and Components," VMware vSAN Administration Guide, 2023. [Online]. Available: https://core.vmware.com/resource/vmware-vsan-design-guide

[7] M. K. Aguilera, W. Chen, and S. Toueg, "Heartbeat: A timeout-free failure detector for quiescent reliable communication," in Proc. 11th International Workshop on Distributed Algorithms, Saarbrucken, Germany, 1997, pp. 126-140.

[8] A. Lopez, "High availability and disaster recovery concepts, design, implementation," BMC Software White Paper, 2022.

[9] D. Patterson et al., "Recovery-oriented computing (ROC): Motivation, definition, techniques, and case studies," UC Berkeley Computer Science Technical Report UCB//CSD-02-1175, 2002.

[10] Veeam Software, "Veeam Availability Suite: Modern Data Protection Built for Virtualization," Veeam Technical Documentation, 2023. [Online]. Available: https://www.veeam.com/documentation-guides-datasheets.html

[11] R. S. Sandhu, E. J. Coyne, H. L. Feinstein, and C. E. Youman, "Role-based access control models," IEEE Computer, vol. 29, no. 2, pp. 38-47, Feb. 1996.

[12] VMware, Inc., "VMware vSphere Security Configuration Guide," VMware vSphere 8.0 Documentation, 2023. [Online]. Available: https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-security/

[13] Center for Internet Security, "CIS VMware ESXi 8.0 Benchmark," CIS Benchmarks, v1.0.0, Oct. 2023.

[14] National Institute of Standards and Technology, "NIST Special Publication 800-125: Guide to Security for Full Virtualization Technologies," U.S. Department of Commerce, Jan. 2011.

[15] A. Gupta and D. Singh, "Storage performance benchmarking: A comprehensive survey," ACM Computing Surveys, vol. 55, no. 3, pp. 1-39, 2022.

[16] Intel Corporation, "IOmeter User Guide," Intel Open Source Technology Center, 2021. [Online]. Available: https://www.iometer.org/doc/

[17] FIO Development Team, "Flexible I/O Tester (FIO) Documentation," Version 3.35, 2023. [Online]. Available: https://fio.readthedocs.io/

[18] M. Ben-Yehuda et al., "The turtles project: Design and implementation of nested virtualization," in Proc. 9th USENIX Symposium on Operating Systems Design and Implementation (OSDI), Vancouver, BC, Canada, 2010, pp. 423-436.

[19] S. A. Weil, S. A. Brandt, E. L. Miller, D. D. E. Long, and C. Maltzahn, "Ceph: A scalable, high-performance distributed file system," in Proc. 7th USENIX Symposium on Operating Systems Design and Implementation (OSDI), Seattle, WA, USA, 2006, pp. 307-320.

[20] D. A. Patterson, G. Gibson, and R. H. Katz, "A case for redundant arrays of inexpensive disks (RAID)," in Proc. ACM SIGMOD International Conference on Management of Data, Chicago, IL, USA, 1988, pp. 109-116.

[21] G. J. Popek and R. P. Goldberg, "Formal requirements for virtualizable third generation architectures," Communications of the ACM, vol. 17, no. 7, pp. 412-421, July 1974.

[22] M. Rosenblum and T. Garfinkel, "Virtual machine monitors: Current technology and future trends," Computer, vol. 38, no. 5, pp. 39-47, May 2005.

[23] C. Clark, K. Fraser, S. Hand, J. G. Hansen, E. Jul, C. Limpach, I. Pratt, and A. Warfield, "Live migration of virtual machines," in Proc. 2nd USENIX Symposium on Networked Systems Design and Implementation (NSDI), Boston, MA, USA, 2005, pp. 273-286.

[24] M. Armbrust et al., "A view of cloud computing," Communications of the ACM, vol. 53, no. 4, pp. 50-58, Apr. 2010.

[25] P. Mell and T. Grance, "The NIST Definition of Cloud Computing," NIST Special Publication 800-145, National Institute of Standards and Technology, U.S. Department of Commerce, Sept. 2011.

[26] S. Rose, O. Borchert, S. Mitchell, and S. Connelly, "Zero Trust Architecture," NIST Special Publication 800-207, National Institute of Standards and Technology, U.S. Department of Commerce, Aug. 2020.

[27] International Organization for Standardization, "ISO/IEC 27001:2022 - Information Security, Cybersecurity and Privacy Protection - Information Security Management Systems - Requirements," ISO, Geneva, Switzerland, 2022.

[28] Sophos Ltd., "The State of Ransomware 2024," Sophos Technical Report, 2024. [Online]. Available: https://www.sophos.com/en-us/whitepaper/state-of-ransomware

[29] Grand View Research, "Hyper-Converged Infrastructure Market Size, Share and Trends Analysis Report, 2024-2030," Grand View Research, Inc., 2024. [Online]. Available: https://www.grandviewresearch.com/industry-analysis/hyper-converged-infrastructure-hci-market

[30] K. Hamid, M. W. Iqbal, M. Aqeel, T. Rana, and M. Arif, "Cyber Security: Analysis for Detection and Removal of Zero-Day Attacks (ZDA)," 2023, pp. 172–196. doi: 10.1201/9781003190301-10.

[31] K. Hamid, M. waseem Iqbal, M. Aqeel, X. Liu, and M. Arif, "Analysis of Techniques for Detection and Removal of Zero-Day Attacks (ZDA)," 2023, pp. 248–262. doi: 10.1007/978-981-99-0272-9_17.

[32] N. Malik et al., Behavior and Characteristics of Ransomware - A Survey. 2024, p. 05. doi: 10.1109/ICCR61006.2024.10532983.

[33] K. Khaliq, N. Rahim, K. Hamid, M. Ibrar, U. Ahmad, and M. Ullah, Ransomware Attacks: Tools and Techniques for Detection. 2024, p. 5. doi: 10.1109/ICCR61006.2024.10532926.

[34] M. W. Iqbal, K. Hamid, M. Ibrar, A. Delshadi, N. Zairah, and Y. Mahmood, "Meta-Analysis and Investigation of Usability Attributes for Evaluating Operating Systems," Migration Letters, vol. 21, pp. 1363–1380, Apr. 2024.

[35] K. Hamid et al., "ML-based Meta-Model Usability Evaluation of Mobile Medical Apps," International Journal of Advanced Computer Science and Applications (IJACSA), vol. 15, no. 1, Art. no. 1, 33/30 2024, doi: 10.14569/IJACSA.2024.0150104.

[36] S. Riaz et al., "Software Development Empowered and Secured by Integrating A DevSecOps Design," Journal of Computing & Biomedical Informatics, p. 02, Mar. 2025, doi: 10.56979/802/2025.

[37] M. Ibrar et al., "Econnoitering Data Protection and Recovery Strategies in the Cyber Environment: A Thematic Analysis," International Journal for Electronic Crime Investigation, vol. 8, Dec. 2024, doi: 10.54692/ijeci.2024.0804216.

[38] M. Nazir et al., "The Silent Guard: ML-Based Zero-Knowledge Proofs in Blockchain Security," vol. 3, pp. 1459–1477, Nov. 2025, doi: 10.5281/zenodo.17615295.

[39] S. Muqadir et al., "ML-BASED Analytical Study on the Level of Security of User Data on Social Media Website," vol. 3, pp. 882–895, Nov. 2025, doi: 10.5281/zenodo.17614882.

[40] A. Tahir, P. K. Hamid, P. D. A. Kahloon, and P. S. Zubair, "Cyber Sovereignty Challenges: A Strategic Framework for National Data Protection Using Blockchain Authentication," Contemporary Journal of Social Science Review, vol. 3, no. 3, pp. 1314–1328, Aug. 2025, doi: 10.63878/cjssr.v3i3.1100.

[41] Y.-S. Tsung, C.-H. Liu, et al., "Performance analysis in HyperFlex and vSAN hyper convergence platforms for online course consideration," IEEE Access, 2022, doi: 10.1109/ACCESS.2022.3224435.

[42] VMware, Inc., "vSphere High Availability Deployment Best Practices," VMware Technical White Paper, 2022. [Online]. Available: https://docs.vmware.com/en/VMware-vSphere/