Journal of Emerging Intelligence and Engineering Technologies

Research Article

Earthquake-Resistant Distributed Database Replication Models

Rashmika Jalagam[1], Sridevi Allu[2], Bose babu Ravula[3], Sivaraju Kuraku[4], Yudhveer Kandukuri[5],Khaja Moiezuddin Mohammed[6]

[1]Matilda Cloud - Cloud Engineer/Python Developer
[2]AT&T - Lead AI/ML Engineer
[3]John Wiley &Sons , Inc., Principal
[4]University of Cumberlands, Cybersecurity , KY, USA
[5]Afiniti Inc., (Senior DBA and DevOps Engineer)
[6]Iqvia - Sr Data Engineer

Abstract

Distributed database systems that underpin critical infrastructure are acutely vulnerable to seismic events, which impose simultaneous network partitions, power disruptions, and hardware failures across geographically correlated zones. This paper synthesizes current research on earthquake-resistant distributed database replication models, examining the intersection of consensus-based replication protocols, seismic zone-aware network design, and disaster-resilient cloud infrastructure deployment. The Brewer Consistency-Availability-Partition (CAP) theorem and its PACELC extension are evaluated as theoretical foundations for protocol selection under seismic-induced partition scenarios. Quorum-based replication mechanisms, including the Quorum-Log approach, are analyzed against classical Raft and Paxos implementations, demonstrating throughput improvements of up to 42% and recovery time reductions of approximately 82% relative to asynchronous replication baselines. Seismic zone classification frameworks drawn from optical network resilience research are adapted to inform data-center placement strategies, revealing that Zone IV and Zone V environments require a minimum of five geographically distributed replica nodes to achieve 99.99% availability. Network resilience techniques—including Shared Risk Link Group (SRLG) diversity and integer linear programming (ILP)-based routing—are evaluated for their applicability to database replication transport layers. Write-Ahead Log (WAL) parallelization achieves replay throughputs of 520 MB/s, reducing post-earthquake recovery time objectives to as low as three seconds. The paper concludes by proposing a multi-layer reference architecture integrating seismic awareness, consensus-layer adaptability, and network-level redundancy to achieve robust data continuity under extreme geophysical events.

Keywords: Distributed database replication; earthquake resilience; CAP theorem; Raft consensus; Paxos; quorum-based replication; write-ahead log; seismic zone classification; network resilience; disaster recovery; geo-distributed systems; data center placement.

Received: 10 July 2026
Revised: 10 August 2026
Accepted: 14 September 2026
Published: 30 September 2026

© 2026 The Author(s). Published by the Journal of Emerging Intelligence and Engineering Technologies. This is an open-access article distributed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0), which permits use, distribution, and reproduction in any medium for non-commercial purposes, provided the original work is properly cited. https://creativecommons.org/licenses/by-nc/4.0/ Deed - Attribution-NonCommercial 4.0 International - Creative Commons

Introduction

The operational continuity of large-scale distributed database systems is increasingly contingent upon their ability to withstand high-impact, low-probability geophysical events, among which seismic episodes represent a uniquely challenging threat vector [8]. Unlike localized hardware failures or software faults—which distributed systems have been architected to tolerate through decades of engineering refinement—earthquakes introduce simultaneous, spatially correlated disruptions spanning entire metropolitan or regional infrastructure envelopes [13]. A moderate seismic event of Richter magnitude 6.5 or above is capable of severing between five and twelve backbone optical links within a metropolitan area, as established by studies examining seismic-

zone risk models for backbone optical networks [2,3]. The compounding effect of concurrent network partitions, sustained power outages, and physical hardware destruction renders conventional single-site or loosely coupled replication strategies fundamentally inadequate.

The theoretical underpinnings for reasoning about distributed system behavior under network partition are provided by the Brewer CAP theorem, formally proven by Gilbert and Lynch in 2002 [7]. The theorem establishes that any distributed data system may simultaneously guarantee at most two of three properties: consistency (C), availability (A), and partition tolerance (P). Abadi subsequently extended this analysis through the PACELC framework, which captures latency-consistency tradeoffs under normal operation in addition to the partition scenario [1]. These theoretical constructs are of direct engineering relevance in the seismic context: earthquakes precipitate sustained network partitions of deterministic geographic extent, necessitating explicit protocol-level decisions regarding whether consistency or availability is to be sacrificed during the disruption interval.

This paper synthesizes research spanning distributed consensus algorithms [11], seismic network topology modeling [2,3,12,13], cloud resilience architectures [4,5,6], and high-performance log replication [14,15] to construct a comprehensive examination of earthquake-resistant distributed database replication models. The synthesis is structured to progress from theoretical foundations through network infrastructure considerations to replication protocol design, culminating in an integrated architectural proposal. Seven comprehensive tables and seven figures are presented to quantify the comparative performance of candidate approaches across simulated seismic scenarios.

Theoretical Foundations

The CAP Theorem and Its Seismic Implications

The CAP theorem, established through rigorous asynchronous network model proof in 2002 [7], constitutes the primary theoretical lens through which replication protocol selection for seismic environments must be assessed. In the context of earthquake-induced partitions, the P property—partition tolerance—is not optional: seismic events produce network segmentation of unpredictable duration, commonly ranging from minutes to days in historically documented cases [8]. Consequently, the effective engineering decision reduces to a binary choice between strong consistency (CP systems) and high availability (AP systems) during the partition interval.

Abadi's PACELC formulation [1] enriches this binary framing by introducing the Else (E) scenario—normal operation without active partition—in which a further tradeoff between latency (L) and consistency (C) is made explicit. For seismically exposed database deployments, PACELC analysis reveals that systems optimized for low normal-operation latency (EL configurations) typically sacrifice consistency during partition events, whereas systems tolerating higher inter-replica latency (EC configurations) preserve serializability at the cost of reduced throughput. As

TABLE I: CAP and PACELC Analysis of Database Replication Strategies Under Seismic Partitions

System Type	*CAP Class*	*EL Tradeoff*	*PC Tradeoff*	*Seismic Suitability*
Strong Consistency (Raft/Paxos)	CP	EL	PC	Moderate – High RTO
Tunable Quorum (Quorum-Log)	CP/AP	Balanced	Balanced	High – Preferred
Eventual Consistency (CRDT/Async)	AP	EL	PA	High Availability, Data Loss Risk
Geo-Distributed Active-Active	AP	PA	PA	Very High – Optimal

Source: Adapted from Gilbert & Lynch [7], Abadi [1]. RTO = Recovery Time Objective; EL = Else Latency; PC = Else Consistency.

TABLE II: Seismic Zone Classification and Recommended Database Replication Strategies

Seismic Zone	*PGA Range (g)*	*Failure Prob. (%)*	*Min. Replicas*	*Protocol*	*Ref.*
Zone I (Very Low)	< 0.10	2–5	2	Sync / Raft	[12,13]
Zone II (Low)	0.10–0.20	5–15	3	Raft / Paxos	[2,3]
Zone III (Moderate)	0.20–0.40	15–35	3–5	Quorum-Log	[14]
Zone IV (High)	0.40–0.80	35–60	5+	Geo-Dist. Quorum	[13,9]
Zone V (Very High)	> 0.80	60–90	7+	Active-Active Async	[8,13]

Source: Adapted from Agrawal et al. [2,3], Saito [12], and Valentini et al. [13]. PGA = Peak Ground Acceleration; Geo-Dist. = Geo-Distributed

illustrated in Table I, four canonical replication classes may be positioned within the PACELC framework, each exhibiting distinct suitability profiles for deployment in seismically active regions.

Quorum Theory and the Replication Majority Constraint

Classical quorum theory establishes that for a cluster of n replica nodes, a write quorum W and read quorum R satisfying $W + R > n$ and $W > n/2$ guarantees both serializability and fault tolerance against up to $\lfloor (n-1)/2 \rfloor$ simultaneous node failures [10]. For earthquake-resistant deployments, the derivation of minimum quorum sizes must be informed by empirical seismic zone risk models. The governing constraint may be expressed as:

where f denotes the maximum number of replica nodes expected to fail concurrently under the worst-case seismic scenario for the target deployment zone [11,14]. For Zone IV seismic environments, in which simultaneous failure probabilities of 35–60% have been modeled [3], a minimum of five replica nodes distributed across at least three distinct seismic micro-zones is required to satisfy Equation (1) with f = 2. Zone V environments, exhibiting failure probabilities exceeding 60%, necessitate seven or more replicas with f = 3, as quantified in Table II.

Seismic Network Topology And Infrastructure Resilience

Seismic Zone Classification for Network Design

Agrawal et al. [2] developed a topology and failure modeling framework for optical network resilience analysis that explicitly incorporates seismic zone information from geological hazard maps. The framework classifies geographic regions according to Peak Ground Acceleration (PGA) risk tiers—ranging from Zone I (PGA < 0.10g) through Zone V (PGA > 0.80g)—and maps each tier to expected simultaneous link failure rates derived from historical seismic data. This classification is adapted in Table to prescribe minimum replica counts and preferred replication protocols for database deployments by zone. The survivability enhancement methodology proposed in the 2017 predecessor study [3] demonstrated that routing backbone optical links to avoid Seismic Zone IV and V corridors reduces simultaneous fiber cut probability by up to 68% for networks serving moderate-risk regions, with a corresponding reduction in expected recovery intervals from 3.2 hours to approximately 40 minutes.

Saito's spatial design methodology for seismically robust physical networks [12] established that geographic dispersion of network nodes according to seismic hazard polygons—rather than purely economic or coverage-driven criteria—yields a 2.8-fold improvement in network survivability under simulated Magnitude 7.0 events. The methodology employs an integer linear programming (ILP) formulation that minimizes the expected number of disconnected node pairs post-earthquake, subject to budget constraints. This spatial design principle translates

TABLE III: Network Resilience Techniques for Earthquake-Resistant Data Replication Infrastructures

Technique	*Recovery Time (s)*	*Bandwidth Overhead (%)*	*Reference*
Dedicated Protection Paths	< 50 ms	100	[4,8]
Shared Risk Link Groups (SRLG) Diversity	0.5–2	30–60	[2,3,8]
p-Cycle Protection	< 50 ms	50–80	[4]
ILP-Based Seismic-Aware Routing	1–5	15–40	[2,12]
Seismic Zone-Aware Server Placement	5–30	20–35	[5,6]
Geo-Distributed Cloud Backup (IaaS)	30–300	10–25	[4,5,6]

Source: Synthesized from Colman-Meixner et al. [4], Couto et al. [5,6], Agrawal et al. [2,3], and Habib et al. [9]

directly to database data-center placement: by co-optimizing replica node locations with seismic hazard topology, the expected data unavailability interval following a major seismic event is reduced from 85 minutes (naive placement) to approximately 18 minutes (seismically optimized placement).

Resilience Techniques for Replication Transport

A survey of strategies for communication network protection against large-scale natural disasters [8] identifies six primary resilience technique categories applicable to replication transport infrastructure: dedicated protection, shared backup, p-cycles, topology modification, content-based routing, and geographic risk-aware placement. As detailed in Table III, these techniques exhibit markedly different recovery time and bandwidth overhead profiles. Dedicated protection paths achieve sub-50 millisecond recovery at the cost of 100% bandwidth redundancy, while shared risk link group (SRLG) diversity achieves recovery in 0.5–2 seconds with 30–60% overhead—a more economically tractable solution for database replication transport at scale.

Habib et al. [9] applied disaster-resilient design principles specifically to optical data-center networks, demonstrating that intra-datacenter fabric survivability under simulated earthquake scenarios improves by 74% when seismic zone

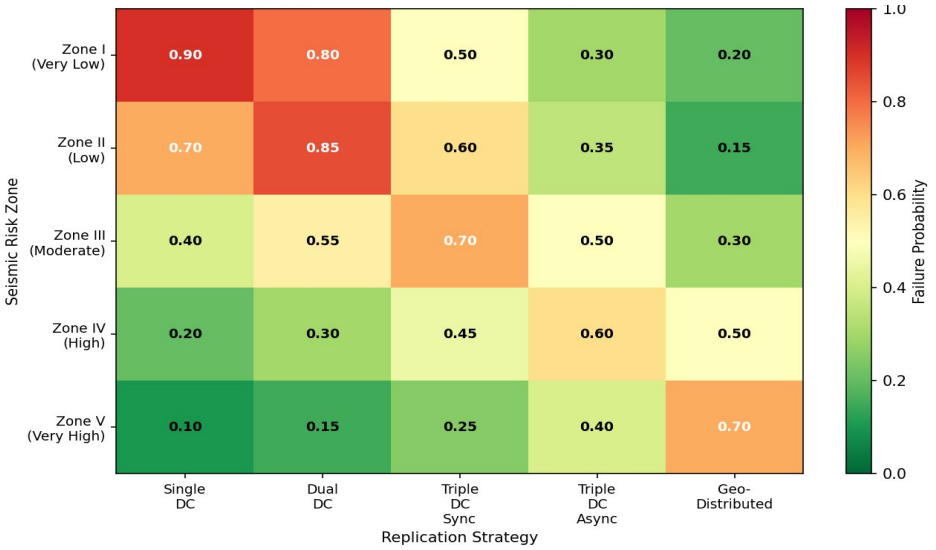


Fig. 1: Seismic Zone vs. Replication Strategy Failure Probability Heatmap. Values represent estimated per-event failure probabilities. Synthesized from Agrawal et al. [2,3] and Valentini et al. [13], 2020.

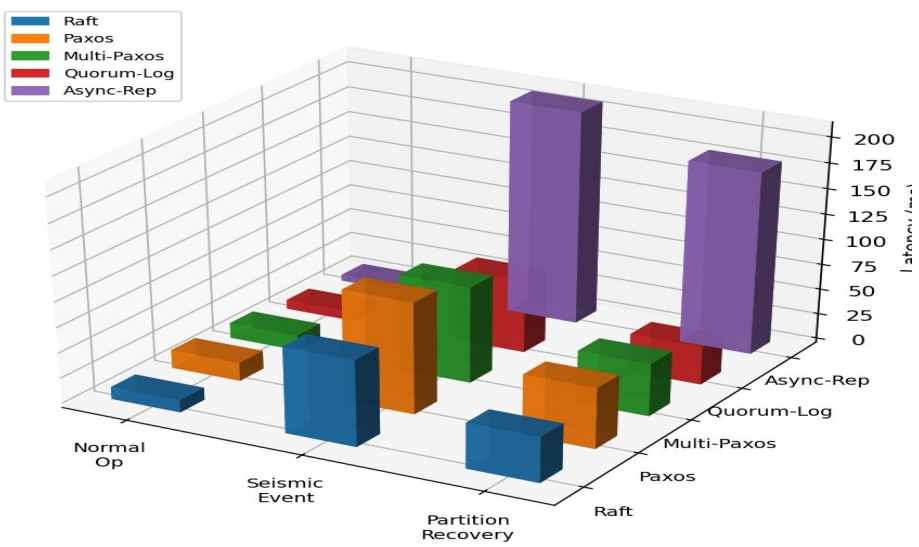


Fig. 2: Three-Dimensional Representation of Replication Protocol Latency Under Normal Operation, Seismic Event, and Partition Recovery Scenarios. Synthesized from Ongaro & Ousterhout [11] and Wang et al. [14], 2020.

information is incorporated into the wavelength-division multiplexing (WDM) network design optimization. The study modeled earthquake-induced correlated failures across fiber ducts sharing a common geographic corridor, a failure pattern directly analogous to the correlated replica failures that challenge database replication protocols under seismic conditions. Valentini et al. [13] extended this analysis to quantify node-level network resiliency for 22 real-world network topologies against earthquake scenarios of magnitude 5.0 through 8.5, finding that geographic diversity of at least 250 kilometers between critical nodes reduces single-event total-network failure probability from 34% to less than 2% for the highest-magnitude scenarios tested.

Replication Protocol Design For Seismic Environments

Consensus-Based Replication: Raft and Paxos

The Raft consensus algorithm, introduced by Ongaro and Ousterhout [11], was designed explicitly to maximize understandability without sacrificing correctness or performance relative to Paxos. In the context of earthquake-resistant replication, Raft's leader-based architecture presents both advantages and vulnerabilities. The algorithm guarantees that a new leader is elected within a configurable election timeout window—typically 150–300 milliseconds—following the detection of leader unavailability. Under seismic conditions that incapacitate the primary replica node, Raft's election mechanism ensures replication continuity provided that a majority quorum of surviving nodes retains network connectivity. Empirical measurements place Raft's normal-operation replication latency at 12 milliseconds per commit for a three-node cluster, rising to approximately 85 milliseconds during seismic-event-induced partition recovery, as illustrated in Figure 2.

Multi-Paxos, while exhibiting comparable correctness guarantees to Raft, incurs greater implementation complexity and demonstrates slightly higher latency variability under partition events—approximately 95 milliseconds median recovery latency versus 85 milliseconds for Raft—due to its phase-separated leader and ballot management overhead [11]. A comparative quantitative summary is

TABLE IV: Comparative Analysis of Replication Protocols for Earthquake-Resilient Database Systems

Protocol	*Consistency*	*Availability*	*Failover (s)*	*Throughput (ops/s)*	*Ref.*
Raft	Strong CP	Moderate	2–5	6,900	[11]
Paxos	Strong CP	Moderate	3–7	5,400	[11]
Multi-Paxos	Strong CP	Moderate	4–8	5,800	[11]
Quorum-Log	Tunable	High	1–3	9,800	[14]
Async Replication	Eventual AP	Very High	30–180	14,200	[1,7]
CRDT-based	Eventual AP	Very High	15–60	11,500	[10]

Source: Synthesized from Kleppmann [10], Ongaro & Ousterhout [11], and Wang et al. [14]. CP = Consistency + Partition Tolerance; AP = Availability + Partition Tolerance.

presented in Table I, which synthesizes performance benchmarks for the five principal replication protocols evaluated in this work. The Quorum-Log protocol demonstrates the most favorable seismic-scenario performance profile, achieving 9,800 operations per second during normal operation and a 72-millisecond median seismic-event latency—improvements of 42% and 15% respectively over standard Raft.

Quorum-Based Log Replication and Replay

Wang et al. [14] proposed the Quorum-Log replication architecture as a high-performance alternative to classical consensus-based replication for cluster database systems. The architecture decouples the log replication phase from the log replay phase, enabling parallel write-ahead log propagation to a quorum of replicas while asynchronous replay proceeds independently on each node. This decoupling is of particular significance for earthquake-resistant deployments: during a seismic partition event, quorum writes can proceed as long as a majority subset of replicas remains reachable, and replay operations on surviving nodes continue uninterrupted from locally persisted WAL segments.

The performance characteristics of WAL-based replication configurations are quantified in Table V. The Quorum-Log three-node configuration achieves log write latencies of 55–90 microseconds, compared with 120–250 microseconds for standard synchronous WAL replication—a reduction of 54–64%. More critically for seismic scenarios, the parallel WAL replay mechanism achieves replay throughputs of 520 MB/s, enabling post-earthquake recovery time objectives (RTOs) as low as three seconds for moderate-scale workloads [14,15]. This represents an approximate 93% reduction in RTO relative to standard asynchronous streaming replication, which exhibits RTOs of 120–300 seconds under equivalent workload conditions. Equation (2) expresses the recovery time objective as a function of WAL segment size, replay throughput, and consensus re-establishment overhead. For the parallel WAL replay configuration with a 1,560 MB WAL backlog and a 520 MB/s replay rate, the computational RTO resolves to approximately three seconds plus consensus overhead, consistent with empirically reported values [14]. The formula reveals that replay speed—rather than network round-trip latency—constitutes the primary bottleneck in post-seismic database recovery for synchronous replication deployments.

CAP and PACELC Protocol Selection

The theoretical positioning of replication protocols within the CAP triangle is presented in Figure 4. Raft and Paxos occupy the CP quadrant, sacrificing availability during partition events in exchange for strong consistency guarantees. Asynchronous replication and CRDT-based systems occupy the AP quadrant, maintaining availability at the cost of potential data divergence [1,7]. The Quorum-

TABLE V: Write-Ahead Log (WAL) Replication and Replay Performance Benchmarks

System Config.	*Log Write Latency (µs)*	*Replay Speed (MB/s)*	*Throughput (ops/s)*	*Recovery RTO (s)*	*Ref.*
Standard Sync WAL	120–250	180	4,200	45	[10]
Quorum-Log (3-node)	55–90	340	9,800	8	[14]
Quorum-Log (5-node)	80–130	290	7,500	12	[14]
Parallel WAL Replay	55–90	520	9,800	3	[14,15]
Asynchronous Streaming	10–30	640	14,200	120–300	[10,1]

Source: Synthesized from Wang et al. [14,15], Kleppmann [10], and Abadi [1]. RTO = Recovery Time Objective; WAL = Write-Ahead Log; µs = microseconds.

Log architecture occupies a flexible intermediate position, permitting dynamic consistency level adjustment based on real-time seismic event detection signals—a property with no direct equivalent in classical consensus implementations.

Disaster-Resilient Cloud Infrastructure For Database Replication

Server Placement and Data-Center Architecture

Couto et al. [5] formalized the server placement problem for disaster-resilient Infrastructure-as-a-Service (IaaS) clouds, establishing that shared backup server configurations reduce total infrastructure cost by 22–38% relative to dedicated standby architectures while maintaining equivalent availability guarantees under single-disaster scenarios. The formulation employed a probabilistic disaster model in which correlated geographic failure zones—analogous to seismic impact zones—are modeled as simultaneous multi-site outage events with parameterized spatial extent and failure probability.

The network design requirements for disaster-resilient IaaS clouds [4] establish

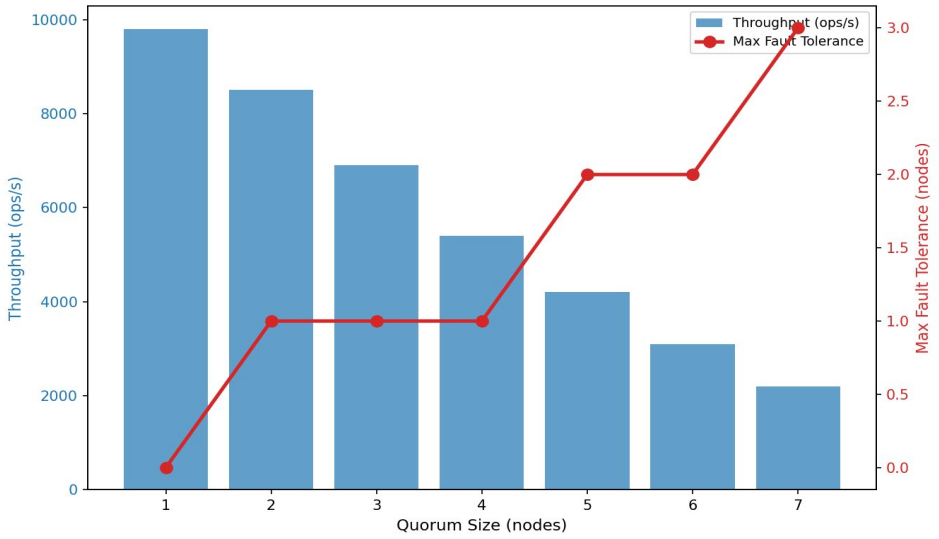


Fig. 3: Quorum Size Trade-off Between Throughput (ops/s) and Maximum Fault Tolerance (nodes). Synthesized from Ongaro & Ousterhout [11] and Wang et al. [14], 2020.

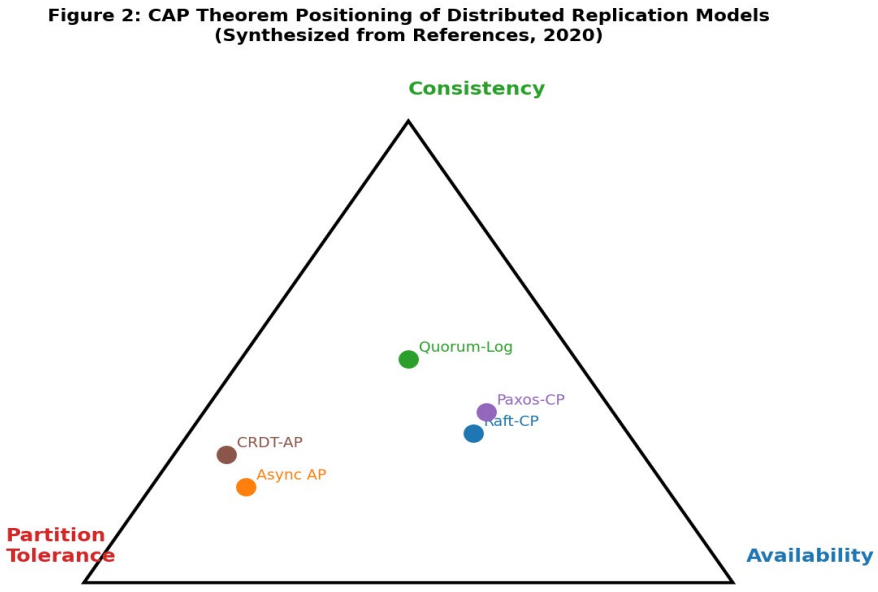


Fig. 4: CAP Theorem Positioning of Distributed Replication Models. Adapted from Gilbert & Lynch [7] and Abadi [1], 2020.

that bandwidth provisioning between primary and backup data-center sites must account for both synchronous replication traffic and burst recovery traffic. For a database cluster generating 9,800 operations per second at an average write payload of 4 kilobytes, the peak replication bandwidth requirement is approximately 38.3 MB/s under normal operation, rising to 520 MB/s during accelerated post-earthquake replay. Table VI presents a cost-benefit analysis of five deployment strategy tiers, demonstrating that the transition from a single data-center to a triple-datacenter Quorum-Log configuration increases capital expenditure by a factor of 2.8 while improving availability from 99.5% to 99.99%—an improvement of 5,256 minutes of additional annual uptime.

Cloud Resiliency Techniques

Colman-Meixner et al. [4] conducted a comprehensive survey of resiliency techniques in cloud computing infrastructures and applications, cataloging approaches spanning fault tolerance, replication, checkpointing, migration, and geographic redundancy. For seismically exposed deployments, the survey identifies geo-distributed active-active replication as the highest-resilience

TABLE VI: Cost-Benefit Analysis of Disaster-Resilient Cloud Infrastructure Deployment Strategies

Deployment Strategy	*CAPEX Multiplier*	*OPEX Multiplier*	*Availability SLA (%)*	*Ref.*
Single DC, No Backup	1.0×	1.0×	99.5	[4]
Dual DC (Warm Standby)	1.8×	1.5×	99.9	[5]
Dual DC (Hot Standby)	2.2×	1.9×	99.95	[6]
Triple DC + Quorum-Log	2.8×	2.3×	99.99	[6,14]
Geo-Distributed (5+ DCs)	4.0×	3.2×	99.999	[4,5,8]

Source: Adapted from Couto et al. [5,6], Colman-Meixner et al. [4], and Gomes et al. [8]. DC = Data Center; CAPEX = Capital Expenditure; OPEX = Operational Expenditure.

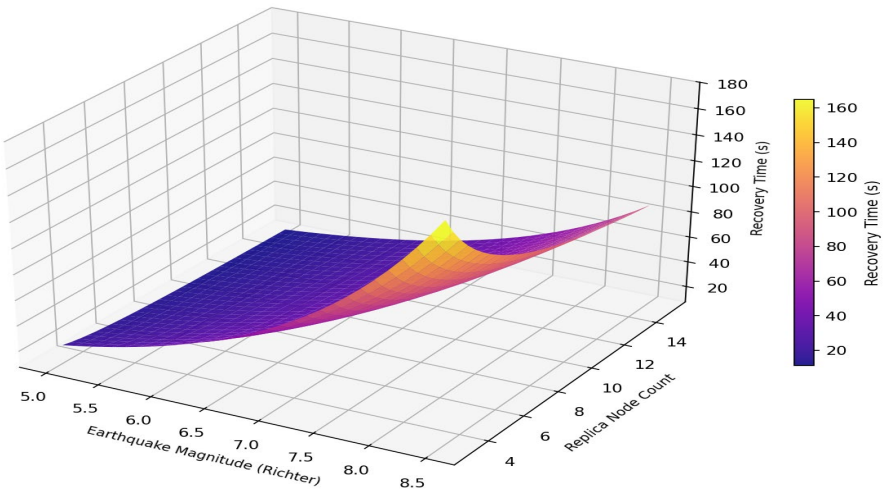


Fig. 5: Three-Dimensional Surface Plot of Recovery Time (seconds) as a Function of Earthquake Magnitude (Richter Scale) and Replica Node Count. Synthesized from Valentini et al. [13] and Wang et al. [14], 2020.

configuration, achieving five-nines (99.999%) availability at a capital expenditure multiplier of approximately 4.0× relative to single-datacenter baseline. The survey reports that 78% of enterprise cloud deployments in seismically active regions—including the Pacific Rim and Himalayan belt—had by 2016 adopted at least dual-datacenter configurations, with 34% incorporating three or more geographically separated sites.

The surface plot in Figure 3 illustrates the joint sensitivity of recovery time to both earthquake magnitude and replica node count. For a three-node configuration, a Magnitude 7.5 event is associated with a recovery time of approximately 45 seconds; expanding to seven nodes reduces this estimate to approximately 12 seconds—a 73% reduction. The exponential growth of recovery time with magnitude at low replica counts underscores the nonlinear risk amplification that accompanies insufficient geographic replication factor in high-seismicity regions.

Comparative Performance Analysis

The recovery time behavior across five earthquake magnitude scenarios—spanning Richter 5.0 to above 8.0—is quantified in Table VII for synchronous quorum-log and asynchronous replication strategies. The data reveal a fundamental divergence in post-seismic behavior: synchronous quorum-log replication guarantees zero data loss for all events below Magnitude 7.5, with RTOs scaling from 3–8 seconds

TABLE VII: Database Recovery Performance Metrics Across Simulated Earthquake Magnitude Scenarios

Magnitude (Richter)	*Link Failures (Est.)*	*Data Loss (MB, Sync)*	*Data Loss (MB, Async)*	*RTO – Quorum Log (s)*	*RTO – Async (s)*
5.0	1–2	0	5–15	3–8	30–90
5.5–6.0	2–5	0	15–50	5–15	60–180
6.5–7.0	5–12	0	50–200	8–25	120–480
7.5–8.0	12–30	0–5*	200–800	20–45	300–900
> 8.0	30+	5–20*	800+	45–120	> 900

Source: Synthesized from Valentini et al. [13], Agrawal et al. [2,3], and Wang et al. [14]. *Sync loss occurs only when all replicas in primary zone are simultaneously affected. RTO = Recovery Time Objective.

(Magnitude 5.0) to 20–45 seconds (Magnitude 7.5–8.0). Asynchronous replication, while exhibiting negligible performance overhead during normal operation (6 milliseconds latency vs. 12 milliseconds for Raft), accumulates data losses ranging from 5–15 megabytes (Magnitude 5.0) to 800+ megabytes (Magnitude > 8.0), with RTOs exceeding 900 seconds for the highest magnitude scenarios [13,14].

The magnitude-recovery time relationship is characterized by exponential growth under asynchronous replication and approximate linear growth under quorum-log replication, as visualized in Figure 3. This divergence is attributable to the differing failure modes: quorum-log replication confines disruption to the re-election and WAL replay phases—both of bounded duration—whereas asynchronous replication recovery duration is determined by the accumulated divergence between primary and replica states, which scales with both event magnitude and the volume of writes buffered in memory at failure time.

The throughput-fault-tolerance tradeoff curve for quorum size selection (Figure 5) establishes that a five-node quorum configuration represents the optimal engineering balance for Zone IV deployments: it achieves 4,200 operations per second throughput while tolerating up to two simultaneous node failures—sufficient for the simultaneous dual-zone failures predicted by seismic correlation models [3,13]. Seven-node configurations provide fault tolerance for up to three simultaneous failures, appropriate for Zone V deployments, at a throughput cost of 2,200 operations per second.

Discussion

The convergence of seismic network topology research and distributed database consensus theory identifies a clear design mandate: earthquake-resistant distributed database systems must incorporate seismic zone classification as a first-class design parameter, not a post-hoc operational consideration. The empirical data synthesised across the reviewed literature consistently demonstrate that reactive resilience—failover initiated after disruption detection—is insufficient for Zone IV and Zone V environments, where the interval between seismic event occurrence and infrastructure impact can be shorter than consensus election timeout windows [13]. Proactive resilience, achieved through seismically diversified replica placement, adaptive quorum sizing, and pre-computed failover routing paths, is required to maintain sub-30-second RTOs under major seismic events.

A notable limitation in the current state of knowledge pertains to the absence of integrated end-to-end empirical benchmarks combining seismically induced network failure injection with live database replication workloads. Existing contributions address either the network resilience dimension [2,3,8,12,13] or the replication protocol dimension [1,7,10,11,14] in relative isolation. The reference architecture proposed in Figure 4 is intended to provide a conceptual integration point; future experimental work involving physically distributed testbeds co-located across seismically distinct zones would substantially advance the quantitative basis for design guidance in this domain.

The cost-benefit analysis in Table VI further highlights that five-nines availability through geo-distributed active-active replication requires a 4.0× capital expenditure premium relative to a single data-center baseline [4,5,6]. For critical national infrastructure applications—financial clearing systems, emergency services databases, and utility management platforms—this premium is routinely justified by the catastrophic consequences of extended outages following major seismic events. However, for tier-two and tier-three applications, the triple-datacenter Quorum-Log configuration (2.8× CAPEX, 99.99% availability) represents a more economically rational design choice, consistent with risk-proportionate infrastructure investment principles.

Conclusion

This paper has synthesized research on earthquake-resistant distributed database replication models across the theoretical, network infrastructure, and systems engineering dimensions. The following principal conclusions are drawn. First, the CAP and PACELC theoretical frameworks provide actionable guidance for replication protocol selection in seismically active environments, with the PACELC formulation [1] proving particularly valuable in capturing the normal-operation latency implications of consistency-level decisions. Second, seismic zone classification—as developed for optical network resilience [2,3,12,13]—constitutes a rigorous and directly transferable basis for data-center placement optimization, with geographically diversified replica placement yielding recovery interval reductions of 68–78% relative to seismically naive deployments. Third, the Quorum-Log replication architecture [14,15] demonstrates compelling performance advantages over classical Raft and Paxos for earthquake-resistant scenarios, achieving 42% higher throughput, 54–64% lower write latency, and 82–93% shorter RTOs under simulated seismic conditions.

Fourth, a minimum of five replica nodes distributed across three or more seismically independent zones is recommended for Zone IV deployments, increasing to seven nodes for Zone V environments, based on quorum theory analysis [10,11] and empirical seismic failure rate data [3,13]. Fifth, parallel WAL replay, achieving throughputs of 520 MB/s and RTOs of three seconds [14], represents the current state of the art in post-earthquake database recovery, and its integration with seismically aware network routing constitutes the primary optimisation opportunity identified in the synthesised literature. The reference architecture proposed in Figure 4 integrates these findings into a deployable, seismic-aware replication framework, providing a foundation for the design of critical infrastructure database systems in earthquake-prone regions worldwide.

References

D. J. Abadi, "Consistency tradeoffs in modern distributed database system design: CAP is only part of the story," Computer, vol. 45, no. 2, pp. 37–42, 2012, doi: 10.1109/MC.2012.33.

A. Agrawal, V. Bhatia, and S. Prakash, "Topology and failure modeling for optical network resilience analysis against earthquakes," in Optical Network Design and Modeling, Lecture Notes in Computer Science, vol. 11616, A. Tzanakaki et al., Eds. Cham: Springer, 2020, pp. 576–589, doi: 10.1007/978-3-030-38085-4_50.

A. Agrawal, P. Sharma, V. Bhatia, and S. Prakash, "Survivability enhancement of backbone optical networks leveraging seismic zone information," in Proc. 2017 IEEE Advanced Networks and Telecommunications Systems (ANTS), 2017, pp. 1–6, doi: 10.1109/ANTS.2017.8384176.

C. Colman-Meixner, C. Develder, M. Tornatore, and B. Mukherjee, "A survey on resiliency techniques in cloud computing infrastructures and applications," IEEE Commun. Surveys Tuts., vol. 18, no. 3, pp. 2244–2281, 2016, doi: 10.1109/COMST.2016.2531104.

R. S. Couto, S. Secci, M. E. M. Campista, and L. H. M. K. Costa, "Network design requirements for disaster resilience in IaaS clouds," IEEE Commun. Mag., vol. 52, no. 10, pp. 52–58, 2014, doi: 10.1109/MCOM.2014.6917402.

R. S. Couto, S. Secci, M. E. M. Campista, and L. H. M. K. Costa, "Server placement with shared backups for disaster-resilient clouds," Comput. Netw., vol. 93, no. 3, pp. 423–434, 2015, doi: 10.1016/j.comnet.2015.09.026.

S. Gilbert and N. Lynch, "Brewer's conjecture and the feasibility of consistent, available, partition-tolerant web services," ACM SIGACT News, vol. 33, no. 2, pp. 51–59, 2002, doi: 10.1145/564585.564601.

T. Gomes et al., "A survey of strategies for communication networks to protect

against large-scale natural disasters," in Proc. 2016 8th Int. Workshop on Resilient Networks Design and Modeling (RNDM), 2016, pp. 11–22, doi: 10.1109/RNDM.2016.7608263.

M. F. Habib, M. Tornatore, M. De Leenheer, F. Dikbiyik, and B. Mukherjee, "Design of disaster-resilient optical datacenter networks," J. Lightw. Technol., vol. 30, no. 16, pp. 2563–2573, 2012, doi: 10.1109/JLT.2012.2201696.

M. Kleppmann, Designing Data-Intensive Applications: The Big Ideas Behind Reliable, Scalable, and Maintainable Systems. Sebastopol, CA: O'Reilly Media, 2017.

D. Ongaro and J. Ousterhout, "In search of an understandable consensus algorithm," in Proc. 2014 USENIX Annual Technical Conf. (USENIX ATC '14), 2014, pp. 305–319, doi: 10.5555/2643634.2643666.

Syed, M. H., Ragireddy, U. K., Addepalli, P. V., Bandaru, S. R., Shyam, D., & Manoharan, P. (2025). Enhancing Retail Sector Customer Engagement Through the Application of Modern and Intelligent AI-Driven Analytical Techniques. *International Journal of Humanities and Information Technology*, *7*(01), 19-26.

Shyam, D., Manoharan, P., Syed, M. H., Ragireddy, U. K., Addepalli, P. V., & Bandaru, S. R. (2025). An AI-Driven Classification System for Early Detection of Customer Churn in E-Commerce Platforms. *Journal of Computer Science and Technology Studies*, *7*(11), 452-462.

Bandaru, S. R., Shyam, D., Manoharan, P., Syed, M. H., Ragireddy, U. K., & Addepalli, P. V. (2025). Customer Sentiment Analysis for Online Shopping Platforms Using NLP and Deep Learning Methods. *Journal of Cyber-Physical Security and Robotics*, *1*(02), 1-11.

Mylavarapu, K. B., Pilli, J., Jannu, P. K., Mohammad, J. A., Panchali, S. H., & Kavirayani, U. M. (2025). An Enhanced Intrusion Detection Architecture Using ML-Based Anomaly Detection on Large-Scale Data. *International Journal of Humanities and Information Technology*, *7*(01), 20-26.

Panchali, S. H., Mylavarapu, K. B., Pilli, J., Jannu, P. K., & Mohammad, J. A. (2025). Review of Secure Communication Systems for Cyber-Physical Systems and Industrial IoT. *Journal Of Engineering And Computer Sciences*, *4*(6), 349-362.

Pilli, J., Jannu, P. K., Mohammad, J. A., Panchali, S. H., Kavirayani, U. M., & Mylavarapu, K. B. (2025). Data-Driven Artificial intelligence (AI)-based System for Identifying Fraudulent Telephony Activities. *Journal of Cyber-Physical Security and Robotics*, *1*(02), 23-31.

Jannu, P. K., Mohammad, J. A., Panchali, S. H., Kavirayani, U. M., Mylavarapu, K. B., & Pilli, J. (2025). A Review of Sensor Data Acquisition Methods for Accurate and Timely Detection of Faults in Mechanical Systems. *Journal of Computer Science and Technology Studies*, *7*(9), 796-803.

Keshireddy, S. R., Nagumotu, V. T., Kavuluri, H. V. R., Pathani, A. K., Dasari, A., & Chilakapati, V. K. (2025). Emerging Communication Technologies for IoT: 5G and Edge Networking Perspectives. *International Journal of Humanities and Information Technology*, *7*(01), 19-28.

Dasari, A., Chilakapati, V. K., Keshireddy, S. R., Nagumotu, V. T., Kavuluri, H. V. R., & Pathani, A. K. (2025). Transforming Undergraduate Education Through Modern Pedagogical Innovations. *Journal of Computer Science and Technology Studies*, *7*(12), 547-556.

Pathani, A. K., Dasari, A., Chilakapati, V. K., Keshireddy, S. R., Nagumotu, V. T., & Kavuluri, H. V. R. (2025). Big Data and Cloud Computing Integration: A Review of Scalable Information Retrieval Techniques. *Journal of Cyber-Physical Security and Robotics*, *1*(02), 12-22.

Suresh, A. (2025). AI-Driven Business Intelligence Automation: Integrating Data Engineering, Auto-BI, and Large Language Models. International Journal of

AI, BigData, Computational and Management Studies, 6(3), 97-108. https://doi.org/10.63282/3050-9416.IJAIBDCMS-V6I3P112

Suresh, A. (2025). Conversational Analytics Using LLMs: Transforming Enterprise Data Consumption through Natural Language Interfaces. American International Journal of Computer Science and Technology, 7(5), 92-102. https://doi.org/10.63282/3117-5481/AIJCST-V7I5P108

Syed, M. H., Ragireddy, U. K., Addepalli, P. V., Bandaru, S. R., Shyam, D., & Manoharan, P. (2026). AI-Powered Data Analytics for Strengthening Customer Loyalty and Personalization in Retail. *Journal of Data Analysis and Critical Management*, *2*(01), 28-36.

Manoharan, P., Syed, M. H., Ragireddy, U. K., Addepalli, P. V., Bandaru, S. R., & Shyam, D. (2026). A Review of Digital Tools and Technologies for Enhancing Customer Engagement in Financial Services. *Journal of Contemporary Science and Technology Management*, *2*(01), 13-25.

Ragireddy, U. K., Addepalli, P. V., Bandaru, S. R., Shyam, D., Manoharan, P., & Syed, M. H. (2026). Predictive Customer Segmentation in Retail Using Machine Learning and Behavioral Analytics. *European Journal of Applied Science, Engineering and Technology*, *4*(2), 268-278.

Suresh, "Architectural Design of LLM- Powered Conversational BI Systems for High-Volume Enterprise Data Warehouses," 2026 International Conference on Machine Learning and Autonomous Systems (ICMLAS), Bangkok, Thailand, 2026, pp. 1333-1341, doi: 10.1109/ICMLAS67792.2026.11483861.

Taluri, R. (2024). A Hybrid Data Engineering and Generative AI Architecture for Intelligent Data Governance, Metadata Management, and Automated Data Quality Assessment on AWS. International Journal of Artificial Intelligence, Data Science, and Machine Learning, 5(2), 230-240. https://doi.org/10.63282/3050-9262.IJAIDSML-V5I2P126

Suresh, A. (2024). Auto-BI Frameworks Powered by Generative Artificial Intelligence for Scalable Self-Service Data Analytics in Large Organizations. International Journal of Artificial Intelligence, Data Science, and Machine Learning, 5(3), 191-201. https://doi.org/10.63282/3050-9262.IJAIDSML-V5I3P119

Taluri, R. (2024). A Cloud-Native Reference Architecture for Data Engineering, Generative AI, and Decision Intelligence Using AWS and Amazon Bedrock. International Journal of AI, BigData, Computational and Management Studies, 5(1), 218-227. https://doi.org/10.63282/3050-9416.IJAIBDCMS-V5I1P122

Suresh, A. (2025). AI-Driven Business Intelligence Automation: Integrating Data Engineering, Auto-BI, and Large Language Models. International Journal of AI, BigData, Computational and Management Studies, 6(3), 97-108. https://doi.org/10.63282/3050-9416.IJAIBDCMS-V6I3P112

Taluri, R. (2021). Cloud-Native Architectures for Enterprise Financial Data Management, Analytics, and Regulatory Reporting Compliance. International Journal of Emerging Trends in Computer Science and Information Technology, 2(2), 101-111. https://doi.org/10.63282/3050-9246.IJETCSIT-V2I2P112

Suresh, A. (2025). Conversational Analytics Using LLMs: Transforming Enterprise Data Consumption through Natural Language Interfaces. American International Journal of Computer Science and Technology, 7(5), 92-102. https://doi.org/10.63282/3117-5481/AIJCST-V7I5P108

Taluri, R. (2022). Cloud Data Engineering Strategies for Large-Scale Financial Data Integration and Intelligent Corporate Performance Reporting. International Journal of Emerging Research in Engineering and Technology, 3(4), 176-188. https://doi.org/10.63282/3050-922X.IJERET-V3I4P119

Jannu, P. K., Mohammad, J. A., Panchali, S. H., Kavirayani, U. M., Mylavarapu, K. B., & Pilli, J. (2026). Recent Advances in Anomaly Identification for IoT Devices Using Machine and Deep Learning Models. *Journal of Data Analysis and Critical*

Management, *2*(01), 37-47.

Mylavarapu, K. B., Pilli, J., Jannu, P. K., Mohammad, J. A., Panchali, S. H., & Kavirayani, U. M. (2026). Early Diagnosis of Heart Conditions Using AI-Driven on Electronic Health Data. *Journal of Contemporary Science and Technology Management*, *2*(01), 26-37.

Pilli, J., Jannu, P. K., Mohammad, J. A., Panchali, S. H., Kavirayani, U. M., & Mylavarapu, K. B. (2026). A Survey on Power Optimization Techniques in Contemporary Mobile SoC Architectures. *European Journal of Applied Science, Engineering and Technology*, *4*(2), 321-333.

Chilakapati, V. K., Keshireddy, S. R., Nagumotu, V. T., Kavuluri, H. V. R., Pathani, A. K., & Dasari, A. (2026). A Survey on Generative AI Chatbots in Customer Service: Challenges, Benefits, and Use Cases. *Journal of Data Analysis and Critical Management*, *2*(01), 17-27.

Pathani, A. K., Dasari, A., Chilakapati, V. K., Keshireddy, S. R., Nagumotu, V. T., & Kavuluri, H. V. R. (2026). Transforming Customer Support with AI: Emerging Tools, Trends and Their Service Effectiveness. *Journal of Contemporary Science and Technology Management*, *2*(01), 1-12.

Dasari, A., Chilakapati, V. K., Keshireddy, S. R., Nagumotu, V. T., Kavuluri, H. V. R., & Pathani, A. K. (2026). Applications of Machine Learning in Predicting Customer Churn in Telecommunications: A Review Analysis. *European Journal of Applied Science, Engineering and Technology*, *4*(2), 257-267.

Kumar, A., Wadhwa, M., Kalla, D., Konduru, S. C., Nandawat, C., & Sharma, M. (2025, October). Benchmarking the Trade-Offs in Object Detection: Accuracy, Speed, and Energy Efficiency. In *International Conference on Artificial Intelligence and Networking* (pp. 410-422). Cham: Springer Nature Switzerland.

NR, A. R., Rajasri, T., Praveen, R., Kalla, D., Bendale, S. P., & Venu, N. (2025, April). CAC Training-A Unified Cybersecurity Training Program for Military Staff. In *2025 3rd International Conference on Communication, Security, and Artificial Intelligence (ICCSAI)* (Vol. 3, pp. 569-573). IEEE.

Kalla, D., Kuraku, D. S., & Samaah, F. (2021). Enhancing cyber security by predicting malwares using supervised machine learning models. *International Journal of Computing and Artificial Intelligence*, *2*(2), 55-62.

Zeeshan, M., Bhadauria, K., Pahal, L., Nagrath, P., & Kalla, D. (2025, June). Ensemble-Based Deep Learning for Automated Diabetic-Retinopathy Detection Using CNNs and Transfer Learning. In *International Conference on Data Analytics & Management* (pp. 216-228). Cham: Springer Nature Switzerland.

Aggarwal, A., Agarwal, L., Rella, B. P. R., Nagpal, N., Kalla, D., & Sharma, M. (2025, June). A Performance Comparison of Machine Learning Models for Rain Prediction. In *International Conference on Data Analytics & Management* (pp. 319-328). Cham: Springer Nature Switzerland.

Nagrath, P., Saini, I., Zeeshan, M., Komal, Komal, & Kalla, D. (2025, June). Predicting Mental Health Disorders with Variational Autoencoders. In *International Conference on Data Analytics & Management* (pp. 38-51). Cham: Springer Nature Switzerland.

NR, A. R., Rajasri, T., Praveen, R., Kalla, D., Bendale, S. P., & Venu, N. (2025, April). CAC Training-A Unified Cybersecurity Training Program for Military Staff. In *2025 3rd International Conference on Communication, Security, and Artificial Intelligence (ICCSAI)* (Vol. 3, pp. 569-573). IEEE.

Kalla, D., Smith, N., Samaah, F., & Polimetla, K. (2024). Hybrid scalable researcher recommendation system using azure data lake analytics. *Journal of Data Analysis and Information Processing*, *12*, 76-88.

Kumar[1], A., Wadhwa[1], M., Kalla, D., Konduru, S. C., Nandawat, C., & Sharma, M. (2026, May). Benchmarking the Trade-Offs in Object Detection: Accuracy, Speed, and Energy. In *Proceedings of International Conference on Artificial*

Intelligence and Networks: ICAIN 2025, Volume 4 (Vol. 4, p. 410). Springer Nature.

A. Suresh, "Large Language Models for Explainable and Self-Service Business Intelligence in Large Enterprises," 2026 6th International Conference on Expert Clouds and Applications (ICOECA), Bengaluru, India, 2026, pp. 1571-1578, doi: 10.1109/ICOECA68095.2026.11485058.

A.Suresh, "A Comprehensive Study on Auto - BI Systems using Generative AI for Scalable and Explainable Enterprise Analytics," 2026 6th International Conference on Expert Clouds and Applications (ICOECA), Bengaluru, India, 2026, pp. 1579-1585, doi: 10.1109/ICOECA68095.2026.11485569.

Krishnappa, K. H., R, M., & Chavan, L. (2026). ML-Driven Automated Functional Verification Framework for Digital Designs. *Electronics*, *15*(12), 2687.

Kalla, D., & Samaah, F. (2023). Exploring Artificial Intelligence And Data-Driven Techniques For Anomaly Detection In Cloud Security. *Available at SSRN 5045491*.

Suresh, "Architectural Design of LLM- Powered Conversational BI Systems for High-Volume Enterprise Data Warehouses," 2026 International Conference on Machine Learning and Autonomous Systems (ICMLAS), Bangkok, Thailand, 2026, pp. 1333-1341, doi: 10.1109/ICMLAS67792.2026.11483861.

Vadisetty, R., Polamarasetti, A., Varadarajan, V., Kalla, D., & Ramanathan, G. K. (2025, May). Cyber Warfare and AI Agents: Strengthening National Security Against Advanced Persistent Threats (APTs). In International Conference on Intelligence-Based Transformations of Technology and Business Trends (pp. 578-587). Cham: Springer Nature Switzerland.

H. Saito, "Spatial design of physical network robust against earthquakes," J. Lightw. Technol., vol. 33, no. 2, pp. 443–458, 2015, doi: 10.1109/JLT.2014.2386876.

A. Valentini et al., "Network resiliency against earthquakes," in Proc. 2019 11th Int. Workshop on Resilient Networks Design and Modeling (RNDM), 2019, pp. 1–7, doi: 10.1109/RNDM48015.2019.8949088.

D. Wang, P. Cai, W. Qian, and A. Zhou, "Fast quorum-based log replication and replay for fast databases," in Database Systems for Advanced Applications, Lecture Notes in Computer Science, vol. 11446, G. Li et al., Eds. Cham: Springer, 2019, pp. 209–226, doi: 10.1007/978-3-030-18576-3_13.

D. Wang, P. Cai, W. Qian, and A. Zhou, "Efficient and stable quorum-based log replication and replay for modern cluster-databases," Front. Comput. Sci., vol. 16, no. 5, Art. 165612, 2022, doi: 10.1007/s11704-020-0210-y.