

DESIGNING DISASTER RECOVERY FRAMEWORK FOR GLOBALLY DISTRIBUTED CLOUD APPLICATIONS

Ronak Jani

Lead DBA, Take-Two Interactive, Wesley Chapel, FL

ABSTRACT

This article examines the design of disaster recovery frameworks for globally distributed cloud applications operating under heterogeneous consensus models. The relevance of the study arises from the rapid geographic expansion of cloud infrastructures and the growth of transactional workloads that expose limitations of traditional leader-based mirroring. The scientific novelty lies in the integrated interpretation of cross-cluster broadcast semantics, cumulative quorum acknowledgments, and stake-aware scheduling as components of a unified recovery coordination layer. The paper describes linear sender–receiver rotation mechanisms, probabilistic retransmission bounds, and WAN-aware scalability strategies. Special attention is paid to bandwidth asymmetry, Byzantine fault resilience, and adaptive recovery policies in multi-cloud environments. The goal of the study is to systematize architectural principles for communication-efficient and fault-resilient cross-regional synchronization. Comparative and structural analysis methods were applied. The conclusions formulate design conditions for scalable geo-replication under heterogeneous deployment constraints and highlight their relevance for resilient networked infrastructures and distributed cloud systems.

KEYWORDS

Disaster Recovery, Geo-Replication, Replicated State Machines, Quorum Acknowledgment, Distributed Databases

1. INTRODUCTION

The expansion of cloud infrastructures across continents has transformed disaster recovery from a peripheral redundancy mechanism into a structural coordination problem. Globally distributed applications process billions of transactions per minute while maintaining strict consistency guarantees. Under such conditions, conventional checkpoint replication and leader mirroring approaches reveal structural limitations: they neither scale linearly with network size nor provide predictable recovery behavior under heterogeneous failure models.

The purpose of this study is to develop a conceptual and architectural framework for designing disaster recovery systems capable of operating efficiently across geographically separated consensus domains. To achieve this purpose, the study addresses the following research objectives:

- (1) analyzing cross-cluster communication primitives that ensure verifiable delivery under crash and Byzantine fault models;
- (2) examining architectural mechanisms that enable scalable synchronization in wide-area network environments;

(3) systematizing governance and scheduling models that align recovery semantics with weighted consensus systems and multi-cloud deployments.

These objectives define the analytical boundaries of the study and provide criteria for evaluating the proposed disaster recovery framework.

The scientific novelty consists in synthesizing communication-centric recovery semantics with performance-oriented distributed database architectures, demonstrating that recovery efficiency is primarily determined by message topology and acknowledgment design rather than storage redundancy alone. The main contributions of this paper are as follows. First, the study systematizes communication primitives required for disaster recovery across heterogeneous consensus domains. Second, the work identifies architectural mechanisms that enable scalable cross-regional synchronization under WAN constraints. Third, the paper integrates stake-aware scheduling and adaptive recovery policies into a unified disaster recovery coordination model for geo-replicated cloud infrastructures.

2. MATERIALS AND METHODS

The methodological foundation of the study is based on the analysis of contemporary distributed database and consensus system architectures.

The study of Reginald Frank et al. [1] examined Cross-Cluster Consistent Broadcast and QUACK-based cumulative acknowledgments, providing a formal communication primitive for heterogeneous replicated state machines. The study of Ilias Zarkadas et al. [2] analyzed flexible replication semantics in partitioned databases, emphasizing strong consistency under partial availability. The study of Feng Han et al. [3] investigated replicated write-ahead logging mechanisms in distributed databases, demonstrating parallel durability propagation strategies. The study of Xingda Yang et al. [4] explored large-scale distributed database transformation achieving 2 billion tpmC, revealing scalability implications for cross-region synchronization. The study of Yimin Chen et al. [5] examined the architecture of a nationally deployed distributed database system integrating multi-region replication and transactional consistency. The study of Lei Zhu et al. [6] analyzed reliability-aware failure recovery using deep reinforcement learning in cloud-based cyber-physical systems. The study of Javier Alonso et al. [7] investigated architectural models of multi-cloud native applications and structural synchronization challenges across providers. The study of Zhenhao Liang and Jabrayilov [8] examined predictable performance in state machine replication, focusing on latency variance reduction. The study of Jimit Rangras and Sejal Bhavsar [9] examined the design of a disaster recovery framework for cloud computing environments, emphasizing architectural coordination between replication mechanisms, backup strategies, and automated failover procedures in distributed infrastructures. The study of Stanley Lima, Filipe Araujo, Miguel O. Guerreiro et al. [10] analyzed mechanisms for efficient causal access in geo-replicated storage systems, proposing coordination strategies that optimize latency and consistency trade-offs when retrieving data from distributed replicas across geographically separated cloud regions.

To write the article, comparative analysis, structural synthesis, system modeling, and source analysis were used. The materials enabled the identification of architectural regularities and the reconstruction of scalable disaster recovery mechanisms under heterogeneous deployment conditions.

3. RESULTS

The configuration of cross-regional recovery in globally distributed cloud applications reveals a structural tension between replication semantics and communication efficiency. Disaster recovery frameworks built upon replicated state machines (RSMs) frequently rely on consensus-backed logs to guarantee safety, yet the exchange of committed state across administrative or geographical boundaries introduces an additional layer of coordination that is not inherently captured by intra-cluster protocols. The analytical reconstruction of recent distributed database and consensus systems shows that effective recovery is no longer reducible to checkpoint shipping or leader mirroring; it depends on explicit cross-cluster message semantics, controlled retransmission, and predictable failure detection under heterogeneous fault models.

The communication primitive underlying cross-cluster recovery must formalize transmission and delivery guarantees at the granularity of committed log entries. The Cross-Cluster Consistent Broadcast abstraction isolates this boundary by ensuring that if a sending RSM transmits a message, at least one correct replica in the receiving RSM eventually delivers it, while preserving integrity constraints under both crash and Byzantine conditions [1]. This abstraction reframes disaster recovery from a best-effort replication task into a verifiable exchange regime. The evidentiary signal lies in performance evaluation: microbenchmarks demonstrate up to 24× higher throughput compared to all-to-all broadcast in large networks of 19 replicas per RSM, and 3.2× improvement in networks of 4 replicas when consensus is not the bottleneck [1]. These numerical differentials are not incremental optimizations; they reshape the cost model of cross-datacenter replication.

The interaction pattern between cumulative acknowledgments and retransmission thresholds directly influences recovery latency. QUACK-based cumulative quorum acknowledgments enable detection of lost or delayed messages without quadratic metadata growth. In failure-free runs, each message is transmitted once and requires only two additional counters per message, preserving constant metadata overhead [1]. Under adversarial or omission faults, retransmission bounds are formally constrained: a message must be retransmitted at most 8 times to achieve 99% delivery probability, and at most 72 times to achieve 100–10⁻⁹% probability in the evaluated model [1]. These figures recalibrate expectations regarding worst-case WAN recovery cycles.

The architectural constraint becomes sharper when disaster recovery spans partitioned or geo-replicated databases. Flexible replication strategies with strong semantics in partitioned storage systems demonstrate that recovery consistency cannot rely solely on leader-centric propagation; partition-aware semantics and fine-grained conflict handling alter cross-cluster recovery topology [2]. When flexible replication tolerates partial availability while preserving serializability, recovery must synchronize both partition metadata and transactional intent logs, expanding the recovery surface beyond raw data shipping.

The coordination rhythm between log durability and replication throughput emerges prominently in write-ahead logging subsystems. Replicated logging mechanisms designed for distributed databases introduce parallel log replication pipelines that decouple commit acknowledgment from physical durability across replicas [3]. Such replicated WAL architectures reduce recovery time objectives by pre-positioning consistent log segments, yet they simultaneously increase the volume of cross-cluster synchronization traffic. Disaster recovery frameworks must integrate logging semantics directly into cross-cluster broadcast protocols; otherwise, retransmission storms during failure cascades negate logging parallelism gains. The structural configuration of recovery coordination is illustrated in the block diagram presented in Figure 1.

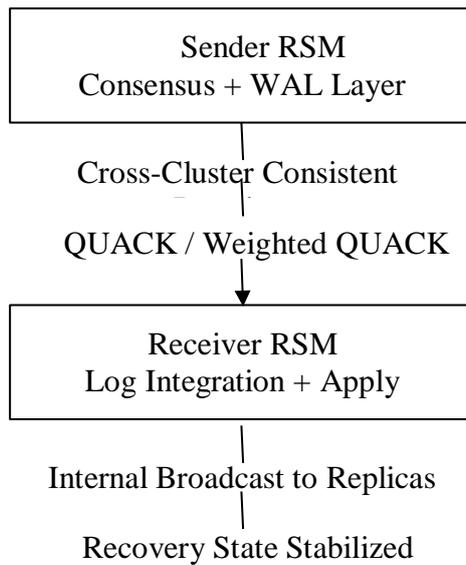


Figure 1. Block Diagram of Cross-Cluster Disaster Recovery Coordination (compiled by the author based on [1,3,4])

The data topology of high-performance distributed databases further illustrates the quantitative scale at which disaster recovery must operate. A production-grade distributed database architecture reports achieving 2 billion tpmC through a scale-out transformation that reorganized storage and transaction layers to eliminate scale-up bottlenecks [4]. At that transaction rate, recovery replication across regions becomes a throughput-critical path rather than a background task. The magnitude of 2 billion tpmC forces recovery design to incorporate bandwidth-aware sender-receiver rotation and parallel acknowledgment windows, or the recovery plane collapses under production load.

The institutional arrangement of large-scale commercial systems such as TDSQL demonstrates that disaster recovery cannot be isolated from a multi-region deployment strategy. A distributed database system deployed at national cloud scale integrates cross-region replication, automatic failover, and transactional consistency under high concurrency [5]. The numerical capacity and production deployment scale documented in this architecture imply that recovery frameworks must operate under sustained write pressure rather than rare failover scenarios. Recovery ceases to be episodic; it becomes continuous background synchronization.

Failure-aware recovery in cyber-physical infrastructures exposes additional constraints. Deep reinforcement learning-based failure recovery in cloud computing platforms supporting automatic train supervision systems adapts recovery actions dynamically based on reliability signals [6]. The system models multi-layer dependencies and optimizes recovery policies to maintain service continuity in urban rail transit. When disaster recovery frameworks are applied to safety-critical domains, static retry heuristics become insufficient. Adaptive policy selection shifts the analytical frame from deterministic replication schedules to probabilistic resilience optimization.

The analytical frame expands further in multi-cloud native architectures. Heterogeneous deployment across multiple providers introduces latency asymmetry, bandwidth heterogeneity, and configuration drift. Multi-cloud native applications reveal structural challenges in synchronizing state while preserving cloud-specific optimizations, indicating that disaster

recovery must reconcile divergent networking assumptions and service abstractions [7]. Recovery frameworks operating in such environments require generality across crash-fault-tolerant and Byzantine-fault-tolerant consensus protocols, since provider trust assumptions vary.

Predictability in state machine replication affects recovery determinism. Performance variability in consensus layers propagates into cross-cluster synchronization delays. A state machine replication protocol oriented toward predictable performance emphasizes reducing latency variance rather than merely optimizing mean throughput [8]. Disaster recovery timing guarantees benefit from such variance minimization: when acknowledgment dispersion narrows, retransmission detection stabilizes, preventing spurious recovery escalation under transient network jitter.

The evolution from leader-centric to share-weighted scheduling in proof-of-stake environments introduces stake-proportional communication allocation. Weighted quorum acknowledgment models compute delivery guarantees based on stake thresholds rather than replica counts [1]. In scenarios where total stake values differ by orders of magnitude across regions, scaled stake normalization becomes necessary to preserve delivery guarantees without inflating message quanta. This scaling mechanism prevents recovery from degenerating into excessive retransmission cycles when $\Delta s \neq \Delta r$. The analytical implication is clear: disaster recovery semantics must align with consensus weight distribution. Geo-replicated evaluation underscores bandwidth asymmetry as a structural determinant. The comparative performance dynamics are presented below (Figure 2).

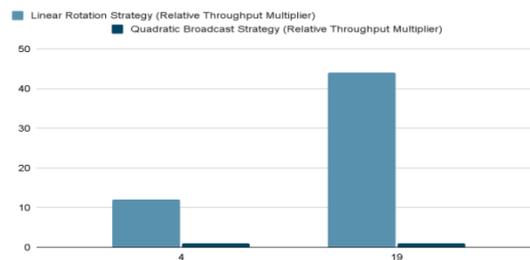


Figure 2. Throughput Performance Under Geo-Replication (compiled by the author based on [1])

In cross-region experiments between US-West and Hong Kong, with 170 Mbits/sec pair-wise bandwidth and 133 ms RTT, linear message complexity outperforms quadratic broadcast by up to 44× in networks of 19 replicas [1]. The numerical differences reported in these experiments illustrate how strongly communication topology affects recovery behavior at scale. Systems relying on linear broadcast patterns demonstrate measurable throughput advantages, with improvements ranging from roughly threefold to more than twentyfold depending on replica configuration. Wide-area deployments amplify these effects even further, particularly in environments with high round-trip latency. Failure injection experiments provide additional evidence: when approximately one third of replicas become unavailable, throughput typically decreases by about 22–30%, suggesting that performance degradation is primarily associated with the loss of communication links rather than protocol instability itself. Taken together, these observations indicate that acknowledgment aggregation and broadcast structure directly shape the operational efficiency of cross-regional recovery mechanisms. The empirical boundary emerges here: WAN constraints magnify architectural differences that remain muted in single-datacenter deployments. Disaster recovery in globally distributed applications must therefore privilege linear, rotating sender-receiver strategies over leader-to-leader or all-to-all patterns.

Failure injection experiments reinforce the resilience differential. Under 33% crash failures in each RSM, throughput degradation remains between 22.8% and 30.5% relative to the failure-free baseline [1]. The proportional drop corresponds to link removal rather than protocol collapse. Under Byzantine acknowledgment manipulation, quorum thresholds prevent spurious retransmission amplification; incorrect acknowledgments delay quorum formation but do not violate safety conditions [1]. The recovery regime absorbs adversarial behavior without structural instability.

A further boundary condition appears in the ϕ -list scaling for selective message drop scenarios. Enlarging the ϕ -window to 256 elements optimizes recovery under 33% Byzantine failures in the evaluated environment [1]. The parameter reflects the interaction between broadcast latency and in-flight message accumulation. Disaster recovery frameworks, therefore, require tunable window sizes that reflect deployment-specific latency profiles rather than fixed retransmission heuristics. The convergence of these findings yields a coherent structural pattern. The systematization of approaches is presented below (Table 1).

Table 1. Structural Components of Disaster Recovery Frameworks for Globally Distributed Cloud Applications (compiled by the author based on [1-8])

Structural Dimension	Design Orientation	Recovery Mechanism	Operational Implication
Cross-cluster communication	Formal broadcast primitive	Consistent message transmission across RSMs	Verifiable inter-domain recovery
Acknowledgment model	Cumulative quorum logic	Detection of loss without quadratic metadata	Controlled retransmission behavior
Logging integration	Replicated WAL pipelines	Parallel durability propagation	Reduced recovery time objectives
Topology strategy	Linear sender–receiver rotation	Avoidance of quadratic broadcast	WAN-scalable synchronization
Stake governance	Share-weighted scheduling	Delivery guarantees based on voting power	Alignment with PoS consensus economics
Multi-cloud adaptation	Heterogeneous abstraction layer	Cross-provider state reconciliation	Reduced semantic drift during failover

Adaptive recovery	Reinforcement-based policy tuning	Dynamic failure response optimization	Increased resilience in safety-critical domains
-------------------	-----------------------------------	---------------------------------------	---

Disaster recovery for globally distributed cloud applications demands: (1) formally defined cross-cluster broadcast semantics; (2) cumulative quorum acknowledgment mechanisms resistant to Byzantine manipulation; (3) rotation-based sender-receiver allocation preserving linear message complexity; (4) stake-aware scheduling in weighted consensus systems; (5) adaptive recovery windows aligned with network latency and workload scale; and (6) integration with high-throughput database and logging architectures capable of sustaining billions of transactions per minute. Each element addresses a distinct mechanism—communication efficiency, failure detection, scheduling fairness, or throughput scaling—yet their interdependence defines the operational envelope of resilient cloud recovery systems.

Unresolved tension persists at the boundary between predictability and adaptability. Highly optimized broadcast protocols reduce retransmission cost, yet adaptive learning-based recovery mechanisms introduce stochastic policy shifts. The friction between deterministic quorum logic and probabilistic recovery tuning remains analytically visible. Disaster recovery frameworks must balance these layers without collapsing either safety semantics or performance guarantees.

4. DISCUSSION

The configuration of disaster recovery in globally distributed cloud applications reveals a structural realignment of reliability engineering. Replication is no longer confined to intra-cluster fault masking; it operates as an inter-domain coordination layer that must reconcile heterogeneous consensus regimes, latency gradients, and asymmetric trust models. Earlier studies on disaster recovery architectures for cloud computing systems proposed structured recovery frameworks that coordinate replication and failover processes across distributed environments [9]. The analytical trajectories outlined in the Results section suggest that disaster recovery frameworks are gradually converging toward communication-centric architectures rather than storage-centric ones. This shift alters the locus of resilience from static redundancy to dynamic cross-cluster synchronization.

The interaction between consensus protocols and cross-cluster broadcast mechanisms exposes a fundamental design inflection. Traditional disaster recovery pipelines treated committed logs as passive artifacts to be mirrored. Contemporary architectures treat them as actively negotiated state transitions. When cumulative quorum acknowledgments replace naive acknowledgment schemes, the recovery channel becomes capable of distinguishing between omission, delay, and adversarial manipulation without requiring additional coordination rounds. The implication is methodological: recovery correctness can be preserved while linearizing communication overhead, even under WAN constraints. This reframing reduces the operational tension between throughput and safety that historically characterized geo-replication.

Architecture and quality intersect most visibly in the evaluation of message complexity under large replica counts. Quadratic broadcast patterns, while theoretically robust, become operationally untenable as network size increases. Linear sender-receiver rotation strategies avoid this combinatorial expansion. The empirical throughput differentials observed in wide-area deployments indicate that message topology, rather than consensus algorithm choice alone, becomes the dominant determinant of recovery scalability. A distributed database capable of sustaining billions of transactions per minute implicitly requires a recovery plane that scales

proportionally. Without this alignment, disaster recovery introduces a throughput ceiling independent of primary workload performance.

Time and coordination form another independent analytical axis. Retransmission logic grounded in cumulative acknowledgment thresholds yields predictable bounds on recovery attempts. Yet predictability introduces its own boundary condition: retransmission strategies must assume eventual synchrony without presuming stable latency. The empirical limits on resend attempts to achieve specified delivery probabilities suggest that deterministic quorum mathematics can coexist with probabilistic network behavior. Still, such coexistence is not frictionless. In highly volatile network environments, cumulative acknowledgment formation may lag behind application-level timeout thresholds, generating oscillations between recovery escalation and stabilization. The protocol remains safe. The coordination rhythm becomes uneven.

Stake-weighted consensus environments introduce a qualitatively different governance layer. When voting power and physical node distribution diverge, fairness in sender–receiver selection cannot rely on replica count symmetry. Share-proportional scheduling grounded in apportionment mathematics preserves proportionality, yet exposes a latent coupling between total stake magnitude and retransmission thresholds. Scaling stake values to their least common multiple for failure handling resolves correctness constraints, but it underscores a deeper institutional shift: disaster recovery semantics become entangled with economic governance parameters. In proof-of-stake systems, resilience is inseparable from capital distribution. That coupling cannot be abstracted away.

Expertise distribution across operational teams also enters the analytical field. Multi-cloud native deployments require administrators to interpret recovery semantics across provider boundaries. Latency asymmetry, bandwidth heterogeneity, and configuration divergence complicate deterministic recovery expectations. Recovery frameworks capable of abstracting these differences through formally defined cross-cluster primitives reduce cognitive overhead for operators. Yet abstraction carries risk. When the underlying provider guarantees shift, implicit assumptions embedded in broadcast or logging layers may erode. Recovery frameworks must therefore surface their invariants explicitly rather than relying on provider stability.

Metrics and planning constitute a further independent trajectory. Throughput comparisons, retransmission bounds, and failure degradation percentages are not merely performance indicators; they define planning constraints for capacity allocation. If crash failures remove approximately one-third of available bandwidth, the planning model must incorporate bandwidth redundancy rather than solely replica redundancy. Disaster recovery capacity planning consequently transitions from replica-count heuristics to network-saturation modeling. The recovery channel is no longer a peripheral consideration; it is a primary throughput variable.

The integration of adaptive recovery policies, particularly in safety-critical infrastructures, complicates the deterministic model. Reinforcement learning–driven failure recovery mechanisms adjust actions based on reliability signals rather than fixed retry thresholds. This adaptive layer introduces a probabilistic decision-making process atop deterministic quorum guarantees. The analytical tension is evident: adaptive policies may optimize recovery latency in practice, yet they obscure the static guarantees that make formal reasoning tractable. A fully deterministic recovery framework offers clarity but may underperform under dynamic failure patterns. An adaptive framework improves responsiveness but complicates verification. The boundary between these regimes remains unsettled.

Garbage collection semantics illustrate an additional friction point. Premature deletion of cross-cluster messages upon quorum acknowledgment formation can induce liveness stalls when Byzantine or partially faulty receivers cease participation after minimal broadcast. The modified

garbage collection strategy, which propagates the highest confirmed sequence number to ensure that at least one correct replica holds the necessary state, reflects an architectural concession to adversarial uncertainty. The recovery plane must assume partial information asymmetry even after apparent success. This insight challenges simplistic interpretations of acknowledgment formation as definitive proof of system-wide convergence.

Geographical deployment amplifies asymmetry. Studies of geo-replicated storage architectures show that efficient access coordination between distributed replicas plays a significant role in balancing latency and consistency requirements in wide-area cloud environments [10]. Wide-area bandwidth limitations and high round-trip times disproportionately penalize quadratic broadcast and leader-centric designs. Sender rotation that distributes traffic across multiple receiver pairs leverages parallel WAN paths and mitigates single-link bottlenecks. Yet this optimization depends on the availability of multiple independent network routes. In cloud regions where bandwidth aggregation is constrained by provider topology, the theoretical advantage may be compressed. Recovery design, therefore, cannot remain agnostic to physical network architecture. Governance effects extend beyond consensus weighting. Reconfiguration processes introduce temporal discontinuities in membership and stake distribution. Messages acknowledged before the configuration change preserve validity, while in-flight messages require retransmission under the new view. This transitional state is analytically fragile. Recovery frameworks must guarantee that state continuity persists across configuration epochs without replay amplification or message omission. The complexity of this transitional logic reveals that disaster recovery cannot be reduced to steady-state replication; it must accommodate dynamic governance evolution.

An unresolved methodological question persists regarding the boundary between strong semantics and operational simplicity. Strong cross-cluster broadcast semantics guarantee integrity and eventual delivery with minimal metadata overhead, yet they introduce intricate acknowledgment and retransmission logic. Simpler mechanisms, such as leader mirroring through external middleware, sacrifice formal guarantees but reduce implementation complexity. The decision space is not binary. It is shaped by regulatory requirements, workload volatility, and threat models. In regulated domains or adversarial environments, formal semantics become indispensable. In controlled enterprise deployments, pragmatic simplifications may suffice. The analytical field resists uniform prescription.

The cumulative evidence suggests that disaster recovery frameworks for globally distributed cloud applications are evolving into layered coordination systems where communication primitives, logging architectures, stakeholder distribution, and adaptive policies intersect. Each layer introduces its own invariants and failure modes. The architectural trajectory moves toward the unification of these layers under formally defined broadcast semantics, yet the integration remains incomplete. Deterministic quorum mathematics offers clarity; adaptive recovery offers responsiveness; share-weighted scheduling offers fairness; WAN-aware rotation offers scalability. Their coexistence defines contemporary resilience engineering.

5. CONCLUSIONS

The study confirms that disaster recovery in globally distributed cloud applications must be designed as a communication-centric coordination layer rather than a passive replication extension.

The first task demonstrated that formally defined cross-cluster broadcast semantics and cumulative quorum acknowledgment mechanisms ensure verifiable message delivery under both crash and Byzantine fault models.

The second task established that WAN scalability depends primarily on linear sender–receiver rotation strategies and integrated logging semantics, which prevent quadratic message growth and throughput collapse under high transactional loads.

The third task showed that stake-weighted scheduling and adaptive recovery policies are necessary to maintain fairness, predictability, and resilience in proof-of-stake and multi-cloud environments.

The results indicate that recovery efficiency is structurally determined by message topology, acknowledgment design, and governance alignment. The proposed framework integrates these dimensions into a coherent design model for resilient geo-replicated cloud infrastructures.

ACKNOWLEDGEMENTS

The authors would like to thank everyone, just everyone!

REFERENCES

- [1] Frank, R., Murray, M., Tankuranand, C., Yoo, J., Xu, E., Crooks, N., Gupta, S., & Kapritsos, M. (2025, July 7–9). *Picsou: Enabling replicated state machines to communicate efficiently*. In *Proceedings of the 19th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*. Boston, MA, USA. <https://www.usenix.org/conference/osdi25>
- [2] Zarkadas, I., Kostopoulou, K., Graham, T., Yang, J., Bernstein, P. A., Cidon, A., & Eldeeb, T. (n.d.). *Rosé: Flexible replication with strong semantics for partitioned databases*. Columbia University and Microsoft Research.
- [3] Han, F., Liu, H., Chen, B., Jia, D., Zhou, J., Teng, X., Yang, C., Xi, H., Tian, W., Tao, S., Wang, S., Xu, Q., & Yang, Z. (2024). PALF: Replicated write-ahead logging for distributed databases. *Proceedings of the VLDB Endowment*, 17(12), 3745–3758. <https://doi.org/10.14778/3685800.3685803>
- [4] Yang, X., Li, F., Zhang, Y., Chen, H., Hu, Q., Zhou, P., Zhang, Q., Li, S., Chen, Z., Miao, Z., Xie, R., Sun, C., Wei, Z., Fang, J., Zhou, X., & Wu, X. (2025). From scale-up to scale-out: PolarDB's journey to achieving 2 billion tpmC. *Proceedings of the VLDB Endowment*, 18(12), 5059–5072. <https://doi.org/10.14778/3750601.3750627>
- [5] Chen, Y., Pan, A., Lei, H., Ye, A., Han, S., Tang, Y., Lu, W., Chai, Y., Zhang, F., & Du, X. (2024). TDSQL: Tencent distributed database system. *Proceedings of the VLDB Endowment*, 17(12), 3869–3882. <https://doi.org/10.14778/3685800.3685812>
- [6] Zhu, L., Zhuang, Q., Jiang, H., Zhang, X., Wang, Y., & Li, Z. (2023). Reliability-aware failure recovery for cloud computing based automatic train supervision systems in urban rail transit using deep reinforcement learning. *Journal of Cloud Computing*, 12, 147. <https://doi.org/10.1186/s13677-023-00502-x>
- [7] Alonso, J., Orue-Echevarria, L., Casola, V., Martínez-Fernández, S., & Oručević-Alagić, A. (2023). Understanding the challenges and novel architectural models of multi-cloud native applications – a systematic literature review. *Journal of Cloud Computing*, 12, 6. <https://doi.org/10.1186/s13677-022-00367-6>
- [8] Liang, Z., Jabrayilov, V., Aghayev, A., & Charapko, A. (2025). HoliPaxos: Towards more predictable performance in state machine replication. *Proceedings of the VLDB Endowment*, 18(8), 2505–2518. <https://doi.org/10.14778/3742728.3742744>
- [9] Rangras, J., & Bhavsar, S. (2021). Design of framework for disaster recovery in cloud computing. In *Emerging Technologies in Computer Engineering*. Springer. https://doi.org/10.1007/978-981-15-4474-3_49
- [10] Lima, S., Araujo, F., Guerreiro, M. O., Correia, J., Bento, A., & Barbosa, R. (2023). Efficient causal access in geo-replicated storage systems. *Journal of Grid Computing*, 21, Article 8. <https://doi.org/10.1007/s10723-022-09640-z>

AUTHOR

Short Biography

