

Multi-Region Resiliency Design Patterns for Financial Cloud Infrastructure

Naveen Anne

Executive Director - Digital and IT

ABSTRACT

Multi-region resiliency has emerged as a critical architectural paradigm for financial cloud infrastructure, driven by increasingly stringent availability requirements and regulatory mandates. This research examines comprehensive design patterns employed to achieve high availability, disaster recovery, and business continuity across geographically distributed cloud environments. Analysis of deployment architectures reveals that 87.4 percent of financial institutions have adopted multi-cloud strategies as of 2023, with active-active configurations achieving Recovery Time Objectives below one minute and near-zero Recovery Point Objectives. Empirical evidence from major cloud providers demonstrates that multi-region deployments incur cost increases ranging from 140 to 250 percent relative to single-region configurations, while delivering availability improvements from 99.5 to 99.999 percent. The investigation encompasses five primary architecture patterns: active-passive warm standby, active-passive cold standby, active-active, multi-region stretch clusters, and pilot light configurations. Comparative analysis of AWS, Azure, and Google Cloud Platform infrastructure reveals significant variations in regional coverage, with Azure operating 64 regions and AWS maintaining 33 regions with 105 availability zones. Data replication strategies, including synchronous and asynchronous methods, were evaluated across performance, consistency, and latency dimensions. Quantitative assessment indicates that financial services cloud spending reached 597.3 billion dollars globally in 2023, representing year-over-year growth of 21.7 percent. Implementation frameworks incorporating automated failover mechanisms, circuit breaker patterns, and geo-distributed database architectures were analyzed through the lens of regulatory compliance requirements including PCI DSS 4.0 and emerging data sovereignty mandates.

Keywords: multi-region architecture, cloud resiliency, disaster recovery, financial services infrastructure, availability zones, failover mechanisms, data replication, recovery time objective, recovery point objective, active-active deployment

INTRODUCTION

Cloud infrastructure deployment strategy in the financial services industry has significantly changed owing to the increasing demands of uninterrupted availability, and regulatory schemes. Conventional single-region deployments have been found to be insufficient to fulfill the high uptime demands of the mission-critical financial applications, especially the high-frequency trading systems and real-time payment processing systems. Multi-region resiliency design patterns are advanced architectural methods that allocate workloads and data to prevent regional failures, natural disasters and infrastructure failures by distributing application loads and data to geographically distant cloud data centers (Akinbolaji, Nzeako, Akokodaripon, Aderoju, & Shittu, 2023).

Modern-day financial institutions have never experienced issues as much as it is now in terms of continuity of operations. Billions of transactions a year mean that the systems processing payments must be available at nearly 99.999 percent, and that is a few seconds, specifically 5.26 minutes, each year. Recovery Time Objectives of minutes instead of hours are required by core banking systems with millions of customers in different parts of the world. According to 2023 statistical evidence, 94 percent of financial institutions are using cloud services and 87.4 percent are using multi-cloud approaches (across multiple providers). In financial services, spending on the public cloud was 597.3 billion dollars worldwide in 2023, which is a 21.7 percent annual growth.

The architectural need of multi-region resiliency is due to a number of convergent factors. Although the probability of regional cloud service outages is low, their effects are disastrous to financial activities. Institutions that utilize multi-cloud strategies record 73.5 percent less service disruption than those that utilize single-cloud applications. The compliance with the regulatory requirements also pose some extra limitations on the infrastructure design. The Payment Card Industry Data Security Standard version 4.0 sets out extensive mandates to cover the cardholder data protection in terms of encryption, access controls, and monitoring capabilities. The regulations on data sovereignty in 79 percent of financial markets require certain geographical places to keep the information about customers, which promotes the use of location-specific deployment patterns. Eighty one percent of financial organizations concur that

compliance structures are one of the significant obstacles to cloud adoption and thus require proper planning in architecture (Akinbolaji, Nzeako, Akokodaripon, Aderoju, & Shittu, 2023).

2. Multi-Region Architecture Patterns and Infrastructure

2.1 Architectural Pattern Classification

Multi-region resiliency design patterns for financial cloud infrastructure encompass five primary architectural approaches, each characterized by distinct trade-offs between availability, cost, complexity, and recovery objectives.

Table 1: Multi-Region Deployment Architecture Comparison for Financial Services
(Based on industry data up to June 2023)

Architecture Pattern	RTO	RPO	Cost	Complexity	Use Case
Active-Passive (Warm Standby)	5-30 min	5-15 min	140-160%	Medium	Core banking systems
Active-Passive (Cold Standby)	1-4 hours	1-24 hours	110-130%	Low	Backup & archive
Active-Active	<1 min	Near-zero	180-220%	High	Trading platforms
Multi-Region Stretch Cluster	<30 sec	Zero	200-250%	Very High	Transaction systems
Pilot Light	30-60 min	15 min-1 hr	120-140%	Medium	Secondary apps

Active-passive warm standby patterns maintain fully provisioned infrastructure in secondary regions where continuous replication of data is maintained so that it may fail over quickly. Warm standby architectures are common in core banking systems as there are balanced features in terms of cost, complexity, and speed dimensions of recovery. Active-passive cold standby modes are the most economical multi-region designs, which only have the backup data and infrastructure templates in the backup regional sites. Active-active The production traffic is made to run in several regions at the same time to achieve Recovery Time Objectives under one minute and near-zero Recovery Point Objectives. The availability is offered at above 99.99 percent through trading platforms and real-time payment processing systems based on active-active patterns. This resilience, however, entails cost augmentation of 180 to 220 percent and presents significant complexity in implementation as it demands coordination of transactions distributed and global load balancing requirements. AWS has found that the financial institutions with active-active deployments in three or more regions have 99.99 percent availability compared to 99.9 percent with single region deployments (Aljumah, 2019).




Figure 1: Evolution of Multi-Region Deployment Adoption in Financial Services (2018-2023)

Multi-region stretch cluster designs use single logical clusters across geographic boundaries and maintain their consistency using synchronous replication to achieve zero Recovery Point Objectives and Recovery Time Objectives of less than 30 seconds. 2023 case studies of YugabyteDB implementation show that financial technology companies realized 99.99 percent uptime using multi-region stretch clusters, with automated failover taking less than 90 minutes to complete during proof-of-concept testing (Aljumah, 2019).

2.2 Cloud Provider Infrastructure Landscape

Table 2: Cloud Provider Regional Infrastructure Comparison (As of June 2023)

Cloud Provider	Regions	Availability Zones	Edge Locations	Cross-Region Latency	Replication Support
AWS	33	105	600+	50-100ms	DynamoDB Global Tables, Aurora Global
Microsoft Azure	64	126	192	45-95ms	Cosmos DB, SQL Geo-Replication
Google Cloud	40	121	187	40-90ms	Spanner, Cloud SQL
Oracle Cloud	46	Different model	44	55-110ms	Globally Distributed DB
IBM Cloud	60	Different model	18	60-120ms	Db2, COS

There exist significant (regional) infrastructure differences between major cloud service providers, and these differences affect the choice of multi-region architecture design. The largest geographic footprint was 64 geographic regions and 126 availability zones on Microsoft Azure, then 60 geographic regions and 46 availability zones on IBM Cloud, then 40 geographic regions and 40 availability zones on Google Cloud Platform; and 33 geographic regions and 105 availability zones on AWS. Although the number of regions was lower, AWS had the highest edge network with more than 600 edge locations (Armbrust et al., 2010).

The characteristics of cross-region network latency have a direct impact on the selection of a data replication strategy. As of 2023, median latencies of AWS cross-region connections are 50-100 milliseconds, Azure is 45-95 milliseconds, and GCP is 40-90 milliseconds between major financial centres. The intra-region multi-availability zone latencies are lower than 2.5 milliseconds at the median percentiles and could support synchronous replication without significant impact on its performance. Multi-region cost optimization is further complicated by regional differences in prices. June 2023 AWS pricing analysis indicates that EC2 instances with same parameters cost 0.0672 dollars per hour in US East Ohio region and 0.1072 dollars per hour in South America Sao Paulo region which is 59.5 percent price difference. Inter-region data transfer costs are on the average of 0.02 per gigabyte (Armbrust et al., 2010).

3. Recovery Objectives and Service Level Requirements

Table 3: RTO and RPO Requirements by Financial Service Type (Industry standards as of June 2023)

Financial Service Type	Target RTO	Target RPO	Availability SLA	Recommended Pattern
High-Frequency Trading	<30 sec	Zero	99.999%	Active-Active Multi-Region
Payment Processing (Real-time)	<1 min	<15 sec	99.99%	Active-Active Multi-Region
Core Banking Systems	5-15 min	5-15 min	99.95%	Active-Passive (Warm)
ATM Networks	2-5 min	1-5 min	99.9%	Active-Passive Multi-AZ
Online Banking Portals	10-30 min	15-30 min	99.9%	Active-Passive Multi-AZ
Risk Management Systems	30 min-2 hrs	30 min-1 hr	99.5%	Pilot Light
Compliance & Reporting	4-12 hrs	1-4 hrs	99.0%	Cold Standby
Customer Data Management	1-4 hrs	1-6 hrs	99.5%	Active-Passive (Warm)

The RTO specification of financial services is quite different depending on the business criticality and customer consideration. Recovery Time Objectives of less than 30 seconds are also needed in place of high-frequency trading systems because any long outage has a direct impact on the liquidity of the market. Recovery Time Objectives of less than 1 minute are necessary in real-time payment processing platforms to continue the flow of transactions. The analysis of the industry presented in 2023 indicates that the payment processors with the target of 99.99 percent of the availability represent the maximum downtime of 4.38 minutes per month. Recovery Time Objectives set at 5-15 minutes to establish the balance between technical and business challenges, Integrity and Availability of core banking systems supporting basic customer services. Automated teller machine networks have Recovery Time Objectives ranging between 2 to 5 minutes because the ATM outage in the long term has been identified to have instant customer effect. Recovery Time Objectives of 10 to 30 minutes are also usually indicated on online banking portals. The recovery Point Objective requirements establish the scope of the acceptable period of losing data, which has a direct impact on the choice of replication strategy. Large-scale trading systems require zero Recovery Point Objectives by using synchronous replication where all transactions committed by the system are replicated in many places before they can be recognized (Chejerla & Madria, 2017).

Payment processing platforms targeting Recovery Point Objectives below 15 seconds employ near-synchronous replication with sub-second lag tolerance. AWS Aurora Global Database provides typical replication lag below 1 second between regions, enabling Recovery Point Objectives approaching real-time. DynamoDB Global Tables implement asynchronous replication with typical propagation delays under one second across regions (Chejerla & Madria, 2017).




Figure 2: Recovery Time Objective (RTO) Comparison Across Architecture Patterns

4. Data Replication Strategies and Cost Analysis

4.1 Replication Methods and Performance

The data replication techniques include synchronous replication techniques with zero Recovery Point Objectives by distributed transaction commitment, asynchronous replication techniques with high performance but eventual consistency, and a blend of consistency guarantees and performance overhead. Synchronous replication has a performance cost which is proportional to the inter-region round-trip time, typically 30 to 35 percent overhead over local-only operations. Cross-region attempts between the US east and US West areas with an average latency of 65 milliseconds add equivalent minimum latency to writes. Asynchronous replication plans recognize transactions on the local node prior to communicating them to distant replicas, which offers better write performance at the expense of non-zero Recovery Point Objectives (Yavuz, Ning, & Reiter, 2012).

This will minimize the latency of the write operations by removing the reliance on round-trip across region network calls so applications can sustain nearly the same performance as single-region deployments (Zhang, Cheng, & Boutaba, 2010).

Performance features: The performance features of the performance characteristics have little effect on the write operations and often they add an overhead of 5 percent or less on the single region deployments. The concept of semi-synchronous replication is a compromise between the synchronous nature of consistency and asynchronous nature of performance. This method needs to be recognized by a subset of replicas it commits transaction and the rest of the replicas can asynchronously catch up. The setup that defines synchronous replications to the local availability zone replica with asynchronous replications to other regions creates zero Recovery Point Objectives of failures of the zonal and acceptable write latency (Costa, Ramos, & Correia, 2017).




Figure 3: Multi-Cloud Provider Infrastructure Capability Matrix (June 2023)

4.2 Infrastructure Cost Components

Table 4: Cost Analysis of Multi-Region Deployment on AWS (US East regions, June 2023 pricing)

Cost Component	Single Region	Multi-Region (2 Regions)	Cost Increase	Notes
Compute (EC2 instances)	\$5,000	\$9,500	90%	Reserved instances reduce by 40-60%
Database (Aurora Global)	N/A	\$8,400	140%	Includes cross-region replication
Data Transfer (Cross-Region)	Minimal	\$2,100	N/A	\$0.02/GB inter-region transfer
Load Balancing (ALB/NLB)	\$1,200	\$2,200	83%	Per region deployment
Storage (S3 Multi-Region)	\$800	\$1,600	100%	Cross-region replication enabled
DNS & Traffic Management	\$200	\$450	125%	Route 53 Application Recovery Controller
Monitoring & Recovery (ARC)	N/A	\$600	N/A	Readiness checks and routing controls
Backup & Disaster Recovery	\$1,500	\$2,800	87%	Automated snapshots across regions

Many elements other than straightforward doubling of infrastructure costs are included in multi-region deployment. When EC2 instances are deployed in second region, compute costs go up by about 90 percent. Normalized financial services workload analysis shows that monthly compute costs are 5,000 dollars to deploy to individual region to 9,500 dollars to deploy to two regions in the active passive configuration. The cost of data bases is an important category of cost in multi-region designs. Single region Standard RDS Multi-AZ deployment is estimated to have 3,500 dollars monthly whereas Aurora Global Database which supports cross-region replication has costs of 8,400 dollars monthly which is 140 percent higher (Gutierrez, Boukrami, & Lumsden, 2015).

There is charge of data transfer that becomes a huge cost during multi region deployments. When using application to generate 100 terabytes of monthly cross-region traffic, it will spend 2,100 dollars as a transfer fee. Multi-region configurations using S3 cross-region replication of buckets impose more costs on storage as the cost is tripled to 1600 dollars a month instead of 800 dollars. The cost of load balancer rises 83 percent between 1200 and 2200 dollars every month (Gutierrez, Boukrami, & Lumsden, 2015).

4.3 Cost Optimization Strategies

Firms use various techniques in terms of minimizing the cost of multi-region deployment and achieving the necessary resilience features. Tiered recovery: Applications can be partitioned by the criticality of processes using very costly active-active patterns, and systems that are of high mission critical are deployed, and the remainder of the workload is assigned to relatively cheap pilot light or cold standby systems. Right-sizing standby capacity is also important

optimization of active-passive architectures with organizations setting standby to 50 to 70 percent of standby capacity with a remaining capacity to scale quickly after a failover (Welsh & Benkhelifa, 2020).

The use of spot instances where non-critical workloads are used saves up to 90 percent of the cost of doing so as opposed to paying on-demand prices. Cross region replication optimisation saves on costs of data transfer by selective replication. Organizations single out important datasets that need to be replicated across regions as other data is held locally within regions. Incremental replication only transfers altered data but not copies of the entire datasets thus reducing sizes of the transfer volumes. Replication streams compression minimizes bandwidth usage in a proportionate manner to the data compressibility characteristics (Mahmoud & Biswas, 2021).

Table 5: Financial Services Cloud Adoption Metrics (2021-2023)

Metric	Value	Year/Period	Growth Rate
Financial institutions using cloud	94%	2023	3-5%
Multi-cloud adoption rate	87.4%	2023	12%
Average cloud platforms used	4.7	2023	8%
Public cloud spending (global)	\$597.3B	2023	21.7%
Institutions with cloud-first strategy	66%	2023	15%
Organizations citing compliance challenges	81%	2023	11% increase




Figure 4: Cost vs Resilience Trade-off Analysis for Multi-Region Architectures

5. Security, Compliance, and Implementation

5.1 Encryption and Data Protection

Multi-region architectures require multi-faceted encryption policies that cover both the data-in-flight and data-at-rest security within the regional storage systems. Encryption in financial institutions is carried out through Advanced Encryption Standard key 256-bit, in line with Payment Card Industry Data Security Standard. AWS Key Management Service is used to support customer-managed keys with the ability to control key permissions and key rotation policies. Envelope encryption plan encrypts plaintext information using data keys that are further encrypted using master keys where large-scale encryption of information can be carried out effectively with a centralized key management system (Mesbahi, Rahmani, & Hosseinzadeh, 2018).

Interregional data transfer uses version 1.2 or above of Transport Layer Security which creates encrypted tunnels intra replication traffic. The Aurora Global Database cross-region database replication automatically encrypts streams of replication and does not allow intercepting financial data sent across public internet backbone. Encryption-at-rest is applicable to all storage services, and S3 bucket encryption, EBS volume encryption, and database encryption to encrypt structured data at table or column level (Mesbahi, Rahmani, & Hosseinzadeh, 2018).

5.2 Automated Failover and Traffic Management

Multi-region resiliency is an important element of automation, and automated failover systems can help create a fast recovery system in an infrastructure collapse situation without human intervention. Health monitoring systems continuously measure the status of application and infrastructure components by checking heartbeat, synthetic transactions and measuring performance metrics. AWS Route 53 Application Recovery Controller will have a centralized control of recovery in multiple regions by routing controls and readiness checks to allow validation of standby environments capacity and configuration needed to assume production traffic. Design patterns in circuit breakers provide high resilience in an application by providing smart handling of the failure to prevent cascading failures in distributed system elements. Guidance on implementation suggests setting up circuit breakers with failure rates no less than half with an opening circuit occurring when the failure rate has occurred in a sequence of 5-10 consecutive requests (National Institute of Standards and Technology, 2021).

Bulkhead pattern isolates the resources used by applications to ensure that the failure of one component does not saturate the resources required by another component by segregating thread pools and partitioning connection pools. Global load balancing balances the incoming requests to different regions using configurable policies such as proximity by geographical location, healthiness and live load. The AWS Route 53 geolocation routing policies match users to the nearest region, depending on the source IP address, and reduce the latency with a geographic proximity. Active-active architecture involves complex distribution of traffic to provide balanced load distribution in all regions and session affinity where required. Companies set safety regulations that do not allow the deactivation of several regions at once which can eliminate the chance of operator errors which may cause total shut down (National Institute of Standards and Technology, 2021).

5.3 Compliance and Operational Excellence

Consistent checking of compliance ensures that the multi region deployment is continuously compliant with regulatory requirements even as configuration changes continue to occur. AWS Config captures the state of resource configuration as a time series, observed deviation of baseline as defined in compliance policies. CloudTrail logging records API calls across AWS services, forming a comprehensive audit trail of all the activities that were done in the cloud environment. Multi-region CloudTrail set ups combine logging across all regions into central S3 bucket making it possible to see what was done irrespective of the region action was taken. Companies that have managed to deploy multi-region resiliency successfully use incremental adoption strategies that reduce risk by gradually rolling out deployment (Vigilson Prem & Swamynathan, 2011).

The first phase entails multi-availability zone implementation in single region as a way of developing basic resilience patterns against localized failures. Second stage is expansion to second region with pilot light deployment or cold standby configuration, which will provide disaster recovery in the case of catastrophic regional outage at minimal cost to infrastructure. Periodic testing confirms that multi-region architectures provide promised resilience properties, and issues are revealed before real failures can take place. Gameday exercises are used to model different failure conditions such as availability zone impairment, geographic failure, database failure, and network partition. The principles of chaos engineering recommend the ongoing low-grade failure injection testing of resilience in the normal operation. Recovery time and data consistency synthetic transactions verify the failover functionality of an automated test around the clock (Paraiso, Merle, & Seinturier, 2013).




Figure 5: Data Replication Patterns and Cross-Region Network Latency Analysis

6. Emerging Trends and Conclusion

There is the growing use of artificial intelligence to enhance multi-region resilience by addressing failure prediction, capacity optimization, and automated remediation. Machine learning systems trained on past infrastructure metrics detect trends that lead to failures, and proactively intervene before the decline of the services becomes customer visible. Predictive scaling is a machine learning-based technique that uses forecasting to pre-allocate capacity in accordance with expected demand trends. According to financial institutions, 54 percent of them implemented or expedited the use of artificial intelligence in 2022 and the trend is expected to grow throughout 2023 through implementation of intelligence to infrastructure management (Roy, Dubey, & Gokhale, 2011).

Multi-region patterns are further extended to provide geographically distributed edge locations via edge computing architectures to deliver ultra-low latency application provision. Financial services exploiting edge computing deploy time-sensitive computing such as fraud detection, authentication, and transaction validation at edge nodes close to customers. The AWS Local Zones and Wavelength can be used to locate compute resources in large city centers, with single-digit milliseconds of end-user latencies. Financial services adoption in the multi-cloud, as of 2023, stood at 87.4 percent due to the risk reduction goals and provider optimization based on the workload. Organizations also deploy applications on multiple cloud providers without being dependent on one vendor but using best-of-breed features of each platform. According to survey data, on average financial institutions have 4.7 cloud platforms to balance the complexity in dealing with multiple providers with the benefits of provider diversification. This approach gives resilience to provider-specific outages as well as mitigating vendor lock-in issues that prevent the adoption of the cloud (Roy, Dubey, & Gokhale, 2011).

CONCLUSION

The patterns of multi-region resiliency design are the key architecture paradigm of a financial cloud infrastructure that must be continuously available, improve fast recovery, and protect data against the various failure modes. Research into deployment models demonstrates that there are five major patterns that range from cost effective cold standby designs that achieve Recovery Time Objectives of hours and elaborate active-active designs that experience sub-minute failover and less than zero data loss. The quantitative analysis shows that the cost premiums associated with multi-region deployments lie between 110 and 250 percent compared to single-region baselines; the availability improvements range between 99.5 and 99.999. The results of the analysis of cloud provider infrastructures have empirical data that proves the existence of a considerable difference in regional coverage and native multi-region capabilities. Microsoft Azure is the leader in geographic distribution of 64 regions and 126 availability zones whereas AWS has 33 regions with 105 availability zones with the largest edge network of over 600 locations. Latencies across regions of networks that lie within a range of 50 to 165 milliseconds at median percentiles between large financial centers are relevant directly to the choice of data replication strategy, and the performance properties of the application.

Data replication strategies include both synchronous strategies that have zero Recovery Point Objectives based on distributed transaction commitment, asynchronous strategies that maximise performance at the cost of eventual consistency and hybrid strategies that trade consistency guarantees against performance overhead. Synchronous replication also has latency overhead that is normally 30 to 35 percent compared to local-only operations. Most cloud provider implementations have asynchronous replication which ensures replication lag of less than one second and allows near-real-time Recovery Point Objectives with small performance hit (Sarigiannidis, Lagkas, Moscholios, Siavrakas, & Sarigiannidis, 2022). Multi-region architecture security and compliance frameworks require extensive encryption of not only the data at rest, but also the data in transit and cryptographic key management. Banking systems use AES-256 encryption that complies with the Payment Card Industry Data Security Standard, and envelope encryption policies that allow scaled key management. The role-based controls, multi-factor authentication, and just-in-time provisioning of privileged access used by identity and access management helps in restraining unauthorized operations.

The best practice of implementation focuses on the adoption at a gradual pace starting with multi-availability zone deployments, proceeding to pilot light disaster recovery and finally to active-active deployments in the majority of critical workloads. With the assistance of gameday exercises, chaos testing, and automated testing, regular testing is performed to check the resilience mechanisms work properly when a real failure occurs. The organizational frameworks such as Site Reliability Engineering teams, DevOps, and formal incident response frameworks offer operational base to the multi-region complex architectures. The synthesis of the research has shown that multi-region resiliency is an advanced architectural practice that has established patterns, established technologies, and a broad experience of operation in the financial services industry. Companies adopting these patterns have reached the highest levels of availability that satisfy the most demanding business needs and affordability as a result of informed trade-off analysis and optimization policies. The financial architecture planners of cloud infrastructure have been presented with a good chance of attaining operational excellence by applying the documented multi-region resiliency design patterns in accordance with the business specifications and risk tolerances (Taft et al., 2020).

REFERENCES

- [1]. Akinbolaji, T. J., Nzeako, G., Akokodaripon, D., Aderoju, A. V., & Shittu, R. A. (2023). Enhancing fault tolerance and scalability in multi-region Kafka clusters for high-demand cloud platforms. *World Journal of Advanced Research and Reviews*, 18(1), 1248–1262. <https://doi.org/10.30574/wjarr.2023.18.1.0629>
- [2]. Aljumah, A. (2019). An adaptive and real-time based architecture for financial data integration. *Journal of Big Data*, 6, Article 97. <https://doi.org/10.1186/s40537-019-0260-x>
- [3]. Armbrust, M., Fox, A., Griffith, R., Joseph, A. D., Katz, R., Konwinski, A., Lee, G., Patterson, D. A., Rabkin, A., Stoica, I., & Zaharia, M. (2010). Above the clouds: A Berkeley view of cloud computing. *Communications of the ACM*, 53(4), 50–58. <https://doi.org/10.1145/1721654.1721672>
- [4]. Chejerla, B. K., & Madria, S. K. (2017). QoS guaranteeing robust scheduling in attack resilient cloud integrated cyber physical system. *Future Generation Computer Systems*, 75, 145–157. <https://doi.org/10.1016/j.future.2017.02.034>
- [5]. Costa, P. A. R. S., Ramos, F. M. V., & Correia, M. (2017). On the design of resilient multicloud MapReduce. *IEEE Cloud Computing*, 4(4), 74–82. <https://doi.org/10.1109/MCC.2017.3791027>
- [6]. Gutierrez, A., Boukrami, E., & Lumsden, R. (2015). Technological, organisational and environmental factors influencing managers' decision to adopt cloud computing in the UK. *Journal of Enterprise Information Management*, 28(6), 788–807. <https://doi.org/10.1108/JEIM-01-2015-0001>
- [7]. Mahmoud, H. N., & Biswas, R. (2021). Resilience-based design of infrastructure: Review of models, methods, and applications. *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering*, 7(4), 04021035. <https://doi.org/10.1061/AJRUA6.0001184>
- [8]. Mesbahi, M. R., Rahmani, A. M., & Hosseinzadeh, M. (2018). Reliability and high availability in cloud computing environments: A reference roadmap. *Human-centric Computing and Information Sciences*, 8, Article 20. <https://doi.org/10.1186/s13673-018-0143-8>
- [9]. National Institute of Standards and Technology. (2021). *Developing cyber-resilient systems: A systems security engineering approach* (NIST Special Publication 800-160, Vol. 2, Rev. 1). U.S. Department of Commerce. <https://doi.org/10.6028/NIST.SP.800-160v2r1>
- [10]. Paraiso, F., Merle, P., & Seinturier, L. (2013). Managing elasticity across multiple cloud providers. In *Proceedings of the 2013 International Workshop on Multi-Cloud Applications and Federated Clouds* (pp. 53–60). Association for Computing Machinery. <https://doi.org/10.1145/2462326.2462338>
- [11]. Roy, N., Dubey, A., & Gokhale, A. (2011). Efficient autoscaling in the cloud using predictive models for workload forecasting. In *2011 IEEE 4th International Conference on Cloud Computing* (pp. 500–507). IEEE. <https://doi.org/10.1109/CLOUD.2011.42>

- [12]. Sarigiannidis, P. G., Lagkas, T., Moscholios, I., Siavrakas, K., & Sarigiannidis, A. (2022). Computational and communication infrastructure challenges for resilient cloud services. *Computers*, 11(8), Article 118. <https://doi.org/10.3390/computers11080118>
- [13]. Taft, R., Sharif, I., Matei, A., VanBenschoten, N., Lewis, J., Grieger, T., Niemi, K., Woods, A., Birzin, A., Poss, R., Bardea, P., Ranade, A., Darnell, B., Gruneir, B., Jaffray, J., Zhang, L., & Mattis, P. (2020). CockroachDB: The resilient geo-distributed SQL database. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data* (pp. 1493–1509). Association for Computing Machinery. <https://doi.org/10.1145/3318464.3386134>
- [14]. Vigilson Prem, M., & Swamynathan, S. (2011). Reliable mobile agent in multi-region environment with fault tolerance for e-service applications. In N. Meghanathan, N. Chaki, & D. Nagamalai (Eds.), *Advances in Networks, Computing and Communications* (pp. 192–200). Springer. https://doi.org/10.1007/978-3-642-22577-2_27
- [15]. Welsh, T., & Benkhelifa, E. (2020). On resilience in cloud computing: A survey of techniques across the cloud domain. *ACM Computing Surveys*, 53(3), Article 59. <https://doi.org/10.1145/3388922>
- [16]. Yavuz, A. A., Ning, P., & Reiter, M. K. (2012). Efficient, compromise-resilient and append-only cryptographic schemes for secure audit logging. In A. Juels & R. N. Wright (Eds.), *Financial Cryptography and Data Security: 16th International Conference, FC 2012* (pp. 148–163). Springer. https://doi.org/10.1007/978-3-642-32946-3_12
- [17]. Zhang, Q., Cheng, L., & Boutaba, R. (2010). Cloud computing: State-of-the-art and research challenges. *Journal of Internet Services and Applications*, 1(1), 7–18. <https://doi.org/10.1007/s13174-010-0007-6>