Disaster Recovery and Business Continuity in Microsoft Azure Cloud Environments

SHEKAR RAO LAKAVATH
Bonneylake, Washington, USA
Independent Researcher
Shekar.lakavath@gmail.com

Abstract

Organizations increasingly depend on continuous access to cloud-hosted applications and data, making disaster recovery (DR) and business continuity (BC) planning a central requirement rather than an optional safeguard. Traditional disaster recovery, built around dedicated secondary data centers, imposes substantial capital and operational costs that are often difficult for small and medium-sized organizations to justify. Cloud computing, and Microsoft Azure specifically, reframes disaster recovery as a pay-as-you-go service built on continuous replication, geographically distributed storage, and automated failover orchestration. This article examines how disaster recovery and business continuity can be implemented for Azure-hosted workloads, disaster-recovery-as-a-service economics, fault tolerance, data deduplication, and empirical cloud outage research, together with two additional sources on project governance and financial infrastructure resilience. The article proposes a disaster recovery and business continuity architecture for Azure spanning primary and secondary region replication, geo-redundant backup, and monitoring-driven failover orchestration, compares six DR/BC techniques against their underlying principles and Azure implementation, and maps common disaster recovery challenges to specific Azure-based mitigations. Two tables and one architecture figure illustrate this analysis. The findings indicate that Azure's native replication and backup services substantially reduce the cost and complexity barriers that limited traditional disaster recovery adoption, but that regular failover testing, consistent security posture across primary and recovery sites, and disciplined recovery-plan documentation remain organizational responsibilities that technology alone does not resolve. The article concludes with practical recommendations for organizations implementing disaster recovery and business continuity on Azure.

Keywords: disaster recovery, business continuity, Microsoft Azure, Azure Site Recovery, high availability, fault tolerance, cloud backup

1. Introduction

Disaster recovery refers to the set of policies, tools, and procedures an organization uses to recover technology infrastructure and systems following a disruptive event, while business continuity refers to the broader organizational capability to maintain essential functions during and after such an event. Historically, achieving robust disaster recovery required organizations to maintain a fully provisioned secondary data center, a capital-intensive approach accessible primarily to large enterprises. Wood, Lagar-Cavilla, Ramakrishnan, Shenoy, and Van der Merwe (2011) demonstrated that cloud computing's pay-as-you-go pricing model could substantially reduce this cost barrier by allowing organizations to provision recovery infrastructure only when needed, rather than maintaining fully staffed and equipped standby facilities continuously.

As organizations have shifted production workloads to Microsoft Azure and other public cloud platforms, disaster recovery has correspondingly shifted from a facilities problem to a data replication, orchestration, and testing problem. Azure's own storage infrastructure, described by Calder et al. (2011) in their account of Windows Azure Storage's architecture for strong consistency and high availability, illustrates how cloud providers build redundancy and durability directly into the underlying platform, providing a foundation that customer-level disaster recovery architecture can build upon rather than duplicate from scratch.

Despite this platform-level redundancy, achieving genuine business continuity still depends on decisions the customer, not the cloud provider, must make: which workloads require cross-region replication, how frequently backups must occur to meet a defined recovery point objective, how failover is orchestrated across interdependent systems, and how regularly recovery procedures are tested. This article addresses three questions. First, what disaster recovery and business continuity principles have been established in the cloud computing literature, and how have they evolved from early cloud-based recovery proposals to more recent empirical reliability research? Second, how can these principles be organized into a coherent disaster recovery and business continuity architecture for Azure? Third, what specific challenges do organizations face implementing disaster recovery on Azure, and how does the platform, combined with sound operational practice, address them? The remainder of the article reviews the relevant literature, proposes a disaster recovery and business continuity architecture for Azure, compares DR/BC techniques against their Azure implementation, maps challenges to mitigations, and discusses practical implications.

2. Literature Review

2.1. Foundations of Cloud Storage Reliability and Redundancy

Early research into cloud storage reliability established the cryptographic and architectural mechanisms needed to guarantee data availability and integrity in multi-tenant, distributed environments. Bowers, Juels, and Oprea (2009) proposed HAIL, a high-availability and integrity layer that allows a set of distributed servers to cryptographically prove to a client that a stored file remains intact and retrievable, formally unifying earlier cryptographic and distributed-systems approaches to storage reliability. Vogels (2009) examined the consistency guarantees underlying large-scale distributed storage systems, arguing that eventual consistency, rather than strong consistency, was often a necessary and acceptable tradeoff for achieving the availability and partition tolerance that distributed cloud storage requires.

Cachin, Keidar, and Shraer (2009) examined the broader trust model implicit in cloud storage adoption, arguing that customers effectively delegate significant control over data availability and integrity to the cloud provider, a delegation that disaster recovery architecture must account for explicitly rather than assume away. Abu-Libdeh, Princehouse, and Weatherspoon (2010) proposed RACS, a redundant array of cloud storage providers that stripes data across multiple independent cloud storage services, demonstrating an early architectural pattern for avoiding single-provider dependency in disaster recovery design.

2.2.Disaster-Recovery-as-a-Service and Economic Models

A distinct line of research examined the specific economics and deployment models of cloud-based disaster recovery. Pokharel, Lee, and Park (2010) proposed an early disaster recovery system architecture built on cloud computing resources, demonstrating that virtualized cloud infrastructure could replicate the functional role of a traditional secondary data center at substantially lower cost. Wood, Lagar-Cavilla, Ramakrishnan, Shenoy, and Van der Merwe (2011) extended this work with PipeCloud, a system using pipelined, causally consistent replication to reduce the speed-of-light latency penalty inherent in geographically distributed disaster recovery, directly addressing the tension between replication consistency and recovery point objectives.

Alhazmi and Malaiya (2012) conducted a comparative assessment of disaster recovery alternatives, evaluating on-site, colocation, and cloud-based approaches against cost and recovery-time criteria, finding that cloud-based disaster recovery offered the most favorable cost-to-recovery-time ratio for organizations without existing secondary-site infrastructure. Alshammari, Alwan, Nordin, and Al-Shaikhli (2017) extended this economic and architectural analysis specifically to single-cloud and multi-cloud disaster recovery deployments, identifying consistency of recovery outcomes across

heterogeneous environments as a persistent challenge distinguishing multi-cloud disaster recovery from single-provider approaches.

2.3. Fault Tolerance and High Availability

Fault tolerance research has examined the mechanisms by which cloud systems detect, isolate, and recover from component failures without requiring a full disaster recovery invocation. Nazari Cheraghlou, Khadem-Zadeh, and Haghparast (2016) surveyed fault tolerance architectures in cloud computing, distinguishing proactive approaches that anticipate and prevent failures from reactive approaches that detect and recover from failures after they occur, and finding that most production systems combine both approaches at different layers of the infrastructure stack. Endo et al. (2016) conducted a systematic review specifically of high availability in cloud environments, cataloguing documented cloud outages and finding that the majority of major, publicly reported outages between 2014 and 2015 stemmed from software or configuration errors rather than physical hardware failure, a finding with direct implications for disaster recovery planning that focuses primarily on hardware redundancy.

Garraghan et al. (2018) extended this empirical focus with an analysis of emergent failures in cloud systems, arguing that cloud reliability at scale must be rethought around failures that emerge from complex interactions between otherwise individually reliable components, rather than assuming failures are attributable to a single identifiable faulty component. Mesbahi, Rahmani, and Hosseinzadeh (2018) proposed a reference roadmap for reliability and high availability in cloud computing environments, organizing the field according to fault detection, fault tolerance, and recovery mechanisms and identifying gaps between academic fault-tolerance research and the reliability mechanisms actually implemented in commercial cloud platforms. Kumari and Kaur (2021) provided a more recent survey of fault tolerance in cloud computing, confirming that replication, checkpointing, and failover remain the dominant fault-tolerance techniques in both research and commercial deployment.

2.4. Data Protection, Backup, and Deduplication

Efficient, cost-effective backup is a prerequisite for practical disaster recovery at scale, motivating substantial research into data deduplication techniques that reduce the storage and bandwidth cost of maintaining redundant data copies. Xia et al. (2016) conducted a comprehensive study of the past, present, and future of data deduplication, tracing the evolution of deduplication techniques from early backup-specific implementations toward general-purpose storage system integration, and identifying the tension between deduplication ratio and the computational overhead of identifying duplicate data

as a persistent design tradeoff. Shin, Koo, and Hur (2017) surveyed secure data deduplication schemes specifically for cloud storage systems, examining how deduplication can be combined with encryption to preserve both storage efficiency and data confidentiality, a combination directly relevant to backup strategies for regulated or sensitive Azure workloads.

Bessani, Correia, Quaresma, André, and Sousa (2013) proposed DepSky, a system for dependable and secure storage across a cloud-of-clouds, demonstrating that combining data across multiple independent cloud providers using secret-sharing and erasure-coding techniques could achieve both availability and confidentiality guarantees exceeding what any single provider offers in isolation, an architectural pattern relevant to organizations pursuing multi-cloud backup strategies alongside Azure.

2.5.Empirical Cloud Outage Research and Security Standards

Empirical studies of real-world cloud failures have provided evidence-based grounding for disaster recovery planning that purely theoretical fault-tolerance research cannot. Gunawi et al. (2014) conducted a large-scale study of more than 3,000 issues reported across widely used open-source cloud system projects, finding that a substantial proportion of the most severe failures resulted from simple mistakes in error-handling code rather than complex distributed-systems edge cases, suggesting that improved software engineering practice, not merely additional redundancy, is an important component of disaster resilience. Kang (2018) examined information system disaster recovery specifically from a business continuity perspective, arguing that recovery time objective and recovery point objective targets should be derived from documented business impact analysis rather than set generically across all systems regardless of their actual criticality.

Jansen and Grance (2011), in NIST Special Publication 800-144, established foundational guidelines for security and privacy in public cloud computing that remain widely referenced in disaster recovery planning, particularly regarding the importance of maintaining consistent security controls across primary and backup environments. Ali, Khan, and Vasilakos (2015) extended this security focus with a broader review of cloud computing security opportunities and challenges, emphasizing that disaster recovery and security cannot be treated as independent concerns, since a compromised recovery environment can itself become a vector for further compromise during an active incident.

Taken together, this literature indicates that effective disaster recovery and business continuity on cloud platforms depends on combining platform-level redundancy mechanisms, of the kind Azure provides natively, with customer-level architectural decisions about replication scope, backup

frequency, failover orchestration, and recovery testing discipline. The following sections apply this combined understanding specifically to Microsoft Azure.

3. A Disaster Recovery and Business Continuity Architecture for Azure

Building on the platform-redundancy and customer-orchestration distinction established in the literature (Calder et al., 2011; Wood, Lagar-Cavilla, Ramakrishnan, Shenoy, & Van der Merwe, 2011), this article proposes a disaster recovery and business continuity architecture for Azure. Figure 1 illustrates the architecture, spanning primary and secondary region replication, geo-redundant backup, and monitoring-driven failover orchestration.

Figure 1. Disaster Recovery and Business Continuity Architecture for Microsoft Azure Cloud Environments

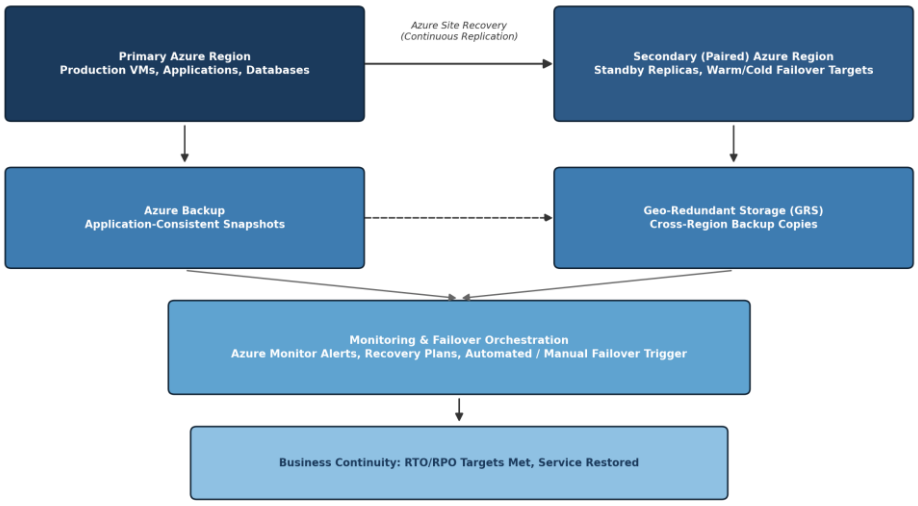

Figure 1. Disaster recovery and business continuity architecture for Microsoft Azure cloud environments, showing continuous replication between primary and secondary regions, geo-redundant backup, monitoring-driven failover orchestration, and the resulting business continuity outcome.

3.1. Regional Replication

Azure Site Recovery continuously replicates production virtual machines and applications from the primary region to a paired secondary region, applying the pipelined replication principles demonstrated by Wood, Lagar-Cavilla, Ramakrishnan, Shenoy, and Van der Merwe (2011) to minimize the gap between recovery point objective targets and replication latency.

3.2. Geo-Redundant Backup

Azure Backup captures application-consistent snapshots, while Geo-Redundant Storage maintains cross-region copies of backup data, applying the deduplication-informed storage efficiency principles documented by Xia et al. (2016) to manage the cost of maintaining these redundant copies at scale.

3.3.Monitoring and Failover Orchestration

Azure Monitor and Azure Service Health provide the early-warning telemetry needed to detect degradation before a full outage occurs, consistent with the emergent-failure detection challenge described by Garraghan et al. (2018). Azure Site Recovery recovery plans encode the specific, scripted sequence of steps required to fail over interdependent systems correctly, addressing the documented-procedure gap that Kang (2018) identifies as essential to meeting business-impact-derived recovery objectives.

3.4.Business Continuity Outcome

The architecture's outcome is measured against defined recovery time objective and recovery point objective targets; meeting these targets, validated through regular test failovers, constitutes the practical definition of business continuity for Azure-hosted workloads.

Table 1. Disaster Recovery and Business Continuity Techniques Compared Against Azure Implementation

DR/BC Technique	Underlying Principle	Azure Implementation
Continuous replication	Maintain a near-real-time copy of production workloads in a secondary location	Azure Site Recovery (VM and application replication to a paired region)
Geo-redundant storage	Store backup and data copies across geographically separated regions	Geo-Redundant Storage (GRS) / Geo-Zone-Redundant Storage (GZRS)
Application-consistent backup	Capture recoverable, transactionally consistent snapshots rather than raw disk copies	Azure Backup with application-consistent snapshot support
Automated failover orchestration	Define and rehearse the sequence of steps required to fail over dependent systems correctly	Azure Site Recovery recovery plans with scripted automation
Continuous monitoring and alerting	Detect degradation or failure early enough to trigger recovery before a full outage	Azure Monitor, Azure Service Health, Azure Advisor
Regular failover testing	Validate that recovery procedures work as intended before a real disaster occurs	Azure Site Recovery test failover (isolated, non-disruptive drills)

4. Comparing DR/BC Techniques and Their Azure Implementation

Table 1 compares six disaster recovery and business continuity techniques established in the literature against their underlying principle and corresponding Azure implementation.

As Table 1 indicates, Azure provides native, largely automated implementations of the core technical mechanisms that the literature identifies as essential to disaster recovery, but regular failover testing, the technique most directly tied to actual recovery success, still depends on organizational discipline rather than platform automation alone, consistent with Kang's (2018) emphasis on documented, rehearsed recovery procedures.

5. Mapping DR/BC Challenges to Azure Mitigations

Table 2 synthesizes the disaster recovery and business continuity challenges most frequently documented in the literature reviewed above and maps each to a specific Azure-based mitigation, along with representative supporting sources.

The mapping in Table 2 indicates that Azure substantially reduces the cost and infrastructure barriers documented in earlier disaster-recovery-as-a-service research, but that organizational challenges, including recovery-plan testing discipline and consistent security posture across primary and recovery environments, persist regardless of the underlying platform, echoing Ali, Khan, and Vasilakos's (2015) argument that security and disaster recovery cannot be addressed independently.

6. Discussion

The analysis presented above supports several practical conclusions for organizations implementing disaster recovery and business continuity on Azure. First, the cost barrier that historically limited disaster recovery to large enterprises has been substantially reduced by cloud-based, pay-as-you-go recovery infrastructure; the economic case that Wood, Lagar-Cavilla, Ramakrishnan, Shenoy, and Van der Merwe (2011) and Alhazmi and Malaiya (2012) established for cloud-based disaster recovery over a decade ago is directly reflected in Azure Site Recovery's consumption-based pricing model, meaning organizations of nearly any size can now justify a documented, tested disaster recovery capability.

Second, empirical research on cloud outages consistently identifies software and configuration errors, not hardware failure, as the dominant cause of significant disruption. Endo et al.'s (2016) and Gunawi

et al.'s (2014) findings both suggest that disaster recovery planning focused primarily on physical redundancy addresses only part of the risk; equal attention should be given to deployment and configuration change management, since a misconfigured change replicated automatically to a secondary region can propagate the same failure that triggered the need for recovery in the first place.

Third, recovery point and recovery time objectives should be derived from documented business impact analysis rather than applied uniformly across all workloads. Kang's (2018) emphasis on business-impact-derived recovery targets implies that Azure customers should classify workloads by criticality and apply differentiated replication frequency and failover automation accordingly, rather than either over-investing in continuous replication for non-critical systems or under-investing in recovery capability for business-critical ones.

Fourth, backup cost management and data protection are not separate concerns from disaster recovery architecture but integral to its sustainability. Xia et al.'s (2016) and Shin, Koo, and Hur's (2017) work on deduplication indicates that the long-term cost-effectiveness of a disaster recovery program depends significantly on efficient backup storage management, particularly as data volumes grow faster than many organizations' backup budgets.

Finally, security and disaster recovery must be planned jointly. Jansen and Grance's (2011) foundational guidance and Ali, Khan, and Vasilakos's (2015) subsequent security review both underscore that a secondary recovery environment inherits the same security requirements as the primary production environment; a disaster recovery architecture that replicates data to a less rigorously secured secondary region has not reduced organizational risk but merely relocated it.

Table 2. Disaster Recovery and Business Continuity Challenges and Azure-Based Mitigations

Disaster Recovery / Business Continuity Challenge	**Azure-Based Mitigation**	**Supporting Literature**
High cost and complexity of maintaining a traditional secondary data center	Pay-as-you-go cloud-based recovery site eliminates the need for a fully provisioned standby facility	Wood, Lagar-Cavilla, Ramakrishnan, Shenoy, & Van der Merwe (2011); Alhazmi & Malaiya (2012)
Single point of failure in storage and data integrity verification	Cryptographic high-availability and integrity layers distribute trust and verification across multiple servers	Bowers, Juels, & Oprea (2009)
Inconsistent recovery outcomes across	Standardized replication and recovery plans applied	Alshammari, Alwan, Nordin, & Al-Shaikhli (2017)

heterogeneous cloud and multi-cloud deployments	consistently across single- and multi-cloud environments	
Long recovery time due to untested or undocumented failover procedures	Scripted, regularly tested recovery plans reduce manual coordination during an actual failover event	Kang (2018); Pokharel, Lee, & Park (2010)
Backup storage costs grow faster than data volume	Data deduplication reduces the volume of redundant data replicated and stored for recovery purposes	Xia et al. (2016); Shin, Koo, & Hur (2017)
Software bugs and configuration errors causing unplanned outages independent of hardware failure	Systematic root-cause analysis of cloud system failures informs more resilient application and infrastructure design	Gunawi et al. (2014); Garraghan et al. (2018)
Difficulty maintaining consistent security and compliance posture across primary and recovery sites	Consistent identity, encryption, and compliance controls applied uniformly to both production and failover environments	Jansen & Grance (2011); Ali, Khan, & Vasilakos (2015)

7. Recommendations

Organizations implementing disaster recovery and business continuity on Azure should: (1) classify workloads by business criticality and derive differentiated recovery point and recovery time objectives from documented business impact analysis, following the approach recommended by Kang (2018); (2) deploy Azure Site Recovery for continuous replication of business-critical workloads to a paired secondary region rather than relying on periodic backup alone; (3) apply Azure Backup with Geo-Redundant Storage for data protection, and evaluate deduplication settings to manage long-term backup storage cost as data volumes grow; (4) conduct regular, scheduled test failovers using Azure Site Recovery's isolated test environment rather than treating recovery plans as a one-time configuration exercise; (5) extend the same identity, encryption, and compliance controls applied in the primary region consistently to the secondary recovery region, following the security guidance of Jansen and Grance (2011); and (6) apply disciplined, data-driven project governance to disaster recovery implementation itself, following the resource allocation and cost-management principles demonstrated by Agumalu (2021) in other operationally complex programs, so that recovery capability is rolled out and tested according to documented business priority rather than implemented inconsistently across the organization.

8. Conclusion

Disaster recovery and business continuity have shifted from a capital-intensive facilities problem to a replication, orchestration, and testing discipline enabled by cloud computing platforms such as Microsoft Azure. This article has reviewed the evolution of disaster recovery research from early cloud storage reliability mechanisms through disaster-recovery-as-a-service economic models, fault tolerance research, data protection techniques, and empirical cloud outage studies, proposed a disaster recovery and business continuity architecture for Azure, and mapped common disaster recovery challenges to specific Azure-based mitigations. The literature reviewed indicates that Azure's native replication, backup, and monitoring services substantially reduce the cost and infrastructure barriers that historically limited disaster recovery adoption, but that recovery-plan testing discipline, workload-appropriate recovery objective setting, and consistent security posture across primary and recovery environments remain organizational responsibilities that the underlying technology, however capable, cannot substitute for. As with other complex operational programs requiring disciplined governance (Adenike, 2018; Agumalu, 2021), realizing the business continuity benefits that Azure's disaster recovery capabilities make possible ultimately depends on well-planned, well-tested, and well-governed implementation. Future research should empirically evaluate recovery time and recovery point outcomes across production Azure disaster recovery deployments to validate the architectural and governance recommendations proposed here.

References

1. Abu-Libdeh, H., Princehouse, L., & Weatherspoon, H. (2010). RACS: A case for cloud storage diversity. In Proceedings of the 1st ACM Symposium on Cloud Computing (SoCC '10) (pp. 229–240). ACM. https://doi.org/10.1145/1807128.1807165
2. Adenike, A. (2018). Strengthening financial market infrastructure in emerging economies: Lessons from central banking operations. SAMRIDDHI: A Journal of Physical Sciences, Engineering and Technology, 10(02), 177–182. https://doi.org/10.18090/samriddhi.v10i02.12
3. Agumalu, S. A. (2021). Project management strategies for improving NGO operational efficiency in Eastern Nigeria: A data-driven assessment of resource allocation, reporting, and project planning practices. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 4(1), 4351–4360.
4. Alhazmi, O. H., & Malaiya, Y. K. (2012). Assessing disaster recovery alternatives: On-site, colocation or cloud. In 2012 IEEE 23rd International Symposium on Software Reliability Engineering Workshops (ISSREW) (pp. 19–20). IEEE. https://doi.org/10.1109/ISSREW.2012.20

5. Ali, M., Khan, S. U., & Vasilakos, A. V. (2015). Security in cloud computing: Opportunities and challenges. Information Sciences, 305, 357–383. https://doi.org/10.1016/j.ins.2015.01.025
6. Alshammari, M. M., Alwan, A. A., Nordin, A., & Al-Shaikhli, I. F. (2017). Disaster recovery in single-cloud and multi-cloud environments: Issues and challenges. In 2017 IEEE 3rd International Conference on Engineering Technologies and Applied Sciences (ICETAS) (pp. 1–5). IEEE. https://doi.org/10.1109/ICETAS.2017.8277910
7. Bessani, A., Correia, M., Quaresma, B., André, F., & Sousa, P. (2013). DepSky: Dependable and secure storage in a cloud-of-clouds. ACM Transactions on Storage, 9(4), Article 12. https://doi.org/10.1145/2535929
8. Bowers, K. D., Juels, A., & Oprea, A. (2009). HAIL: A high-availability and integrity layer for cloud storage. In Proceedings of the 16th ACM Conference on Computer and Communications Security (CCS '09) (pp. 187–198). ACM. https://doi.org/10.1145/1653662.1653686
9. Cachin, C., Keidar, I., & Shraer, A. (2009). Trusting the cloud. ACM SIGACT News, 40(2), 81–86. https://doi.org/10.1145/1556154.1556173
10. Calder, B., Wang, J., Ogus, A., Nilakantan, N., Skjolsvold, A., McKelvie, S., Xu, Y., Srivastav, S., Wu, J., Simitci, H., Haridas, J., Uddaraju, C., Khatri, H., Edwards, A., Bedekar, V., Mainali, S., Abbasi, R., Agarwal, A., Fahim ul Haq, M., … Rigas, L. (2011). Windows Azure Storage: A highly available cloud storage service with strong consistency. In Proceedings of the 23rd ACM Symposium on Operating Systems Principles (SOSP '11) (pp. 143–157). ACM. https://doi.org/10.1145/2043556.2043571
11. Chang, V. (2015). Towards a big data system disaster recovery in a private cloud. Ad Hoc Networks, 35, 65–82. https://doi.org/10.1016/j.adhoc.2015.07.012
12. Endo, P. T., Gonçalves, G. E., Rosendo, D., Gomes, D., Santos, G. L., Moreira, A. L. C., Kelner, J., Sadok, D., & Mahloo, M. (2016). High availability in clouds: Systematic review and research challenges. Journal of Cloud Computing, 5, Article 16. https://doi.org/10.1186/s13677-016-0066-8
13. Garraghan, P., Yang, R., Wen, Z., Romanovsky, A., Xu, J., Buyya, R., & Ranjan, R. (2018). Emergent failures: Rethinking cloud reliability at scale. IEEE Cloud Computing, 5(5), 12–21. https://doi.org/10.1109/MCC.2018.053711662
14. Gunawi, H. S., Hao, M., Leesatapornwongsa, T., Patana-anake, T., Do, T., Adityatama, J., Eliazar, K. J., Laksono, A., Lukman, J. F., Martin, V., & Satria, A. D. (2014). What bugs live in the cloud? A study of 3000+ issues in cloud systems. In Proceedings of the ACM Symposium on Cloud Computing (SoCC '14) (pp. 1–14). ACM. https://doi.org/10.1145/2670979.2670986
15. Jansen, W., & Grance, T. (2011). Guidelines on security and privacy in public cloud computing (NIST Special Publication 800-144). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-144

16. Kang, H.-S. (2018). A study in information system and disaster recovery system for business continuity. Journal of Security Engineering, 15(5), 319–332. https://doi.org/10.14257/jse.2018.10.04

17. Kumari, P., & Kaur, P. (2021). A survey of fault tolerance in cloud computing. Journal of King Saud University - Computer and Information Sciences, 33(10), 1159–1176. https://doi.org/10.1016/j.jksuci.2018.09.021

18. 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

19. Nazari Cheraghlou, M., Khadem-Zadeh, A., & Haghparast, M. (2016). A survey of fault tolerance architecture in cloud computing. Journal of Network and Computer Applications, 61, 81–92. https://doi.org/10.1016/j.jnca.2015.10.004

20. Pokharel, M., Lee, S., & Park, J. S. (2010). Disaster recovery for system architecture using cloud computing. In 2010 10th IEEE/IPSJ International Symposium on Applications and the Internet (SAINT) (pp. 304–307). IEEE. https://doi.org/10.1109/SAINT.2010.62

21. Shin, Y., Koo, D., & Hur, J. (2017). A survey of secure data deduplication schemes for cloud storage systems. ACM Computing Surveys, 49(4), Article 74. https://doi.org/10.1145/3017428

22. Vogels, W. (2009). Eventually consistent. Communications of the ACM, 52(1), 40–44. https://doi.org/10.1145/1435417.1435432

23. Wood, T., Lagar-Cavilla, H. A., Ramakrishnan, K. K., Shenoy, P., & Van der Merwe, J. (2011). PipeCloud: Using causality to overcome speed-of-light delays in cloud-based disaster recovery. In Proceedings of the 2nd ACM Symposium on Cloud Computing (SoCC '11) (Article 17). ACM. https://doi.org/10.1145/2038916.2038933

24. Xia, W., Jiang, H., Feng, D., Douglis, F., Shilane, P., Hua, Y., Fu, M., Zhang, Y., & Zhou, Y. (2016). A comprehensive study of the past, present, and future of data deduplication. Proceedings of the IEEE, 104(9), 1681–1710. https://doi.org/10.1109/JPROC.2016.2571298