Journal of Harbin Engineering University
ISSN: 1006-7043

Vol 46 No. 3
March 2025

Designing High-Availability and Disaster Recovery Solutions for Enterprise Cloud Applications

Harish Kumar. Krishnamurthy. Sukumar

Associate Principal – Cloud Engineering at LTIMINDTREE, Texas, USA

Email – Harish5875@gmail.com

Abstract

In the current digital transformation landscape, it is the availability and reliability of cloud-based business applications that determine the success of an organization. High availability (HA) and disaster recovery (DR) solutions are crucial bricks in minimizing downtime, protecting data integrity, and maintaining continuity of operations during the time of disruption. This document outlines the fundamental concepts and best practices for developing solid HA and DR solutions that are tailored for cloud environments. It addresses the key subjects such as redundancy, fault tolerance, scalability, and automation, and explains how cloud service models (IaaS, PaaS, and SaaS) affect HA/DR designs. Best practices, costs, security and compliance concerns, and testing and verification are also addressed in the article. Case studies are utilized to derive real-world lessons on how to implement and optimize HA and DR strategies for cloud applications.

Keywords—Cloud Computing, High Availability (HA), Disaster Recovery (DR), Business Continuity, Data Protection, Cloud Infrastructure, Scalability, Multi-cloud Strategies, IT Disaster Recovery, Self-Healing Disaster Recovery.

I. INTRODUCTION

Cloud computing is critical to business, offering scalable and affordable infrastructure for mission-critical applications. With businesses going to the cloud, high availability (HA) and robust disaster recovery (DR) solutions must be offered to ensure continuity. HA and DR are vital for planning IT, allowing access, protection of data, and rapid recovery. High availability minimizes downtime through redundancy and failover. Disaster recovery keeps disruption impact at a minimum by restoring services and data rapidly. As businesses rely increasingly on cloud services, it becomes more difficult to guarantee HA and DR. With cloud infrastructure dispersed geographically, enterprises are susceptible to network downtime and hardware failure. While cloud offers scalability and cost savings, providing HA, DR, performance, data consistency, and cost-effectiveness is difficult. HA and DR in cloud apps are topics this article will discuss, underscore business continuity needs, and cite design difficulties for designing a good solution. Let's find out about striking performance, availability, and disaster recovery balance in strong cloud environments. .

II. UNDERSTANDING HIGH AVAILABILITY AND DISASTER RECOVERY

Within the purview of enterprise cloud applications, High Availability (HA) and Disaster Recovery (DR) are two important concepts that provide the reliability, continuity, and resilience of services and data [1]. Though both are essential to reducing disruptions and protecting against failures, they address different facets of sustaining business operations during adverse occurrences and subsequent actions. To be able to deploy HA and DR approaches successfully, organizations must be clear about their definitions, objectives, and underlying strategies.

A. High Availability (HA) Conceptualization

High Availability (HA) is the deliberate design and operation of systems, processes, and infrastructure with the intent to sustain the ongoing availability of a service or application during hardware failures, software bugs, or other forms of interruptions. HA's aim is to prevent or minimize downtime and thus render important services continuously to users, illustrated in Figure 2. High Availability's main aim is to reduce service interruption and allow users uninterrupted access to applications regardless of any failure below. With a high availability system, the failure of one component should not impede the overall system from functioning successfully, with imperceptible user experience. High level design of high availability of (HA) Microsoft azure is shown in figure 1.

Essential Strategies to Attain High Availability Attaining high availability requires the implementation of different techniques and strategies to provide an uninterrupted service even when there is a failure. Some of the most well-known approaches to sustaining high availability include:

1) Failover: Failover systems are designed to switch over to a standby system automatically when the primary system fails. The service is continued by diverting traffic to a replacement instance or server without any manual intervention. Failover systems are commonly used in on-premises and cloud infrastructures.
2) Load Balancing: Load balancing is the process of distributing incoming requests between multiple servers or instances. The more practice here, the better – it ensures that no one server is a single point of failure and that no server becomes overloaded with a greater volume of traffic than it can support. Additionally, it allows requests to be routed to any available resource for better performance and availability.
3) Redundant Resources: Redundancy is the duplication of multiple copies of critical resources, including databases, application servers, or storage devices, in more than one location. This configuration ensures that if one resource is unavailable, a replacement can take its place. Typically, redundant resources are duplicated in various availability zones or geographic locations to protect against localized outages.
4) Geographic Distribution: Services may be replicated across regions or across data centers within the cloud in order to avoid a point of failure. Geographic distribution further reinforces HA by protecting against regional failure, for instance, from natural disasters or power failures.

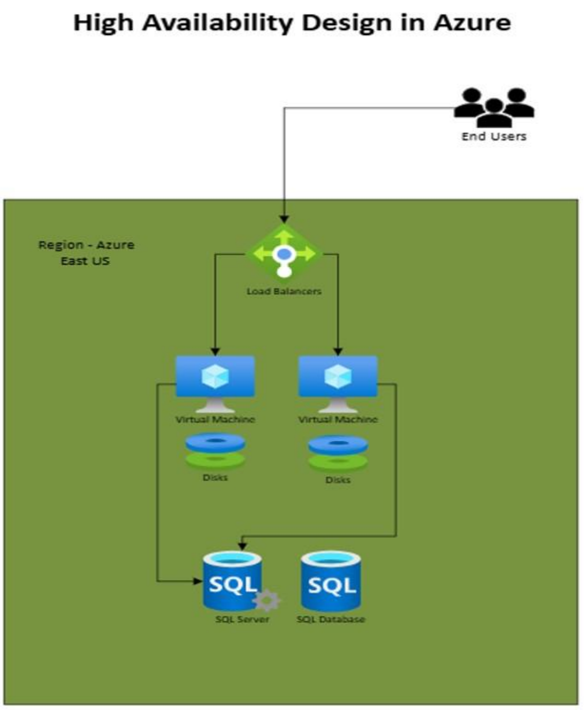

Fig. 1. High availability through VMs and databases across datacenters in same region.

B. Conceptualization of Disaster Recovery (DR)

Disaster Recovery (DR) is the processes, methods, and technological innovation deployed to restore applications, information, and services in the event of a significant disruption, possibly due to system failure, cyber-attack, or natural disaster. DR facilitates the quick recovery of organizational functions with minimal interruption should there be a disaster, hence averting downtime, data loss, and money loss. [2]While High Availability (HA) focuses on being continuously available, DR focuses on recovery from a disastrous failure. DR solutions give organizations the ability to bring their systems back online to an operational state, preferably within an acceptable time, reducing the effect of interruptions to business processes. High level design of high availability of (HA) Microsoft azure is shown in figure 2.

C. Key Approaches to Disaster Recovery

Key Approaches to Disaster Recovery usually include –

1) Data Backup and Replication: The systematic backup of vital information, along with data replication techniques, ensures the successful restoration of data when loss or damage occurs.

The backup information is typically placed in locations that are geographically distant to prevent data loss from localized catastrophes.

2) Automated Recovery Mechanisms: The process of recovery is automated, ensuring system failover and data restoration occur rapidly and properly. Automated systems reduce the potential for human error and improve the rate of recovery, especially in high-pressure environments.

3) Recovery Sites: Other organizations create secondary disaster recovery facilities, as indicated in Figure 1, that are classified as cold, warm, or hot sites, where data and systems are duplicated or pre-configured to support failover in the event of an outage. An example is that a hot site can take over operational responsibilities almost instantly, while a cold site may require extra setup time.

III. RECOVERY TIME OBJECTIVE (RTO) VS. RECOVERY POINT OBJECTIVE (RPO)

There are two important metrics used in disaster recovery planning, namely the Recovery Time Objective (RTO) and the Recovery Point Objective (RPO). These measures enable organizations to specify their acceptable levels of system unavailability and data loss in the event of a disaster:

RECOVERY TIME OBJECTIVE (RTO)

RTO refers to the longest time period for which an application, service, or system can remain unavailable following an event before it begins to have severe adverse effects on the organization. That is, it is the optimum time within which services must be restored in order not to cause significant disruptions of business functions. A lower RTO demands that the recovery be quicker and can involve utilizing cutting-edge or standby infrastructure to achieve this objective.

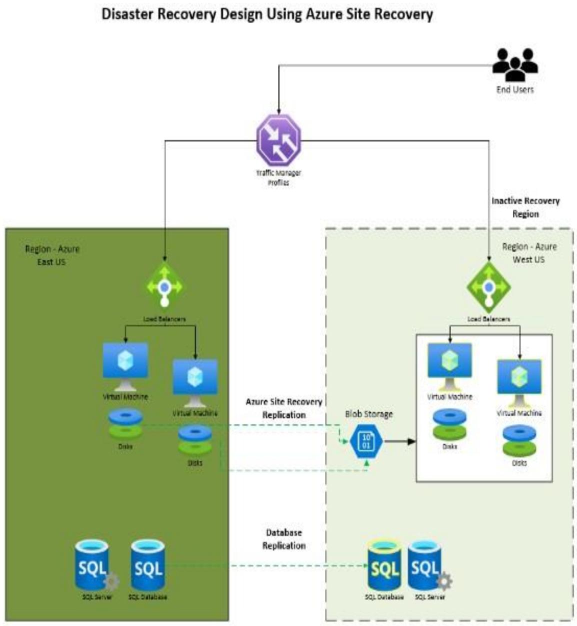

Fig. 2. Disaster Recovery Design by Implementation of Azure Site Recovery in Microsoft Azure..

RECOVERY POINT OBJECTIVE (RPO)

RPO is the amount of acceptable data loss in terms of time. It is the timeframe limit to which the data would be recovered to ensure the organization faces no negative implications. For example, if we have an RPO of four hours, then it requires us to recover the data to an extent of a maximum of four hours before disaster hit [3]. A lower RPO requires more frequent replication or backups so that the amount of loss in case of failure is minimal. Both RTO and RPO assist an organization in determining their disaster recovery objectives and order of priority in the recovery process, trading off recovery speed, cost, and tolerable data loss.

IV. KEY PRINCIPLES OF HIGH-AVAILABILITY (HA) AND DISASTER RECOVERY (DR) DESIGN

A. Redundancy and Fault Tolerance

a) Importance of Redundancy Across Hardware, Data Centers, and Networks: Redundancy is a primary mechanism for maintaining the operability of systems in the event of hardware or network faults. [4] Through the duplication of essential components, the risk of single points of failure (SPOF) is reduced, improving both fault tolerance and overall system resilience.

1) Hardware Redundancy: Involves using multiple servers, storage devices, and networking devices such that in case one fails, another can take over the function without interruption of service.
2) Data Center Redundancy: Means having backup data centers, typically in another geographic location, to provide business continuity in the event a data center is unavailable because of power loss or natural disasters.
3) Network Redundancy: Having several paths within the network and internet connections so that if one fails, traffic is diverted through a different path to offer continuous service.

B. Data Replication and Geographical Distribution

1) Multi-Availability Zone and Multi-Region Disaster Recovery Strategies Spreading infrastructure over multiple regions or availability zones (AZs) introduces additional layers of redundancy [5]. This ensures that if a particular region or AZ encounters a detrimental event, such as a natural disaster or power outage, another region can seamlessly take over operations.
2) Multi-Region Strategies These involve distributing workloads and data across separate physical locations, often spanning different cities or countries, to mitigate interruptions at the regional level.
3) Multi-Availability Zone (Multi-AZ) Strategies These replicate infrastructure across multiple independent data centers within the same region, ensuring high availability and fault tolerance.
4) Synchronous Replication

a) Ensures real-time data replication across multiple sites.

b) Any write to the primary storage is immediately written to secondary storage, maintaining identical copies.

c) Provides maximum data loss protection but may

introduce latency, especially over long distances.

5) Asynchronous Replication

a) Data is transmitted to a backup system with a delay.

b) Improves performance, particularly for long-distance replication, by reducing latency.

c) Introduces a risk of data loss if a failure occurs before replication is completed.

d) Typically used in applications where immediate consistency is less critical or where network bandwidth and latency are limiting factors.

V. AUTOMATION AND DISASTER RECOVERY WORKFLOWS

Automation is defined as the use of scripts, tools, or orchestration platforms to automate the management of failover and recovery procedures [6]. With automation, you completely remove human error and speed up the recovery process during a disaster.

1) Failover automation: Can detect when a system or application is down and automatically reroute traffic to a healthy instance, often across different data centers or regions.
2) Disaster Recovery Workflows: Include predefined steps that can be automatically executed in case of an incident, such as bringing backup systems online, redirecting traffic, restoring backups, or switching to standby systems.

A. Benefits of Infrastructure as Code (IaC) in Automating and Managing HA/DR

Infrastructure as Code (IaC) refers to the process of creating and configuring infrastructure by using machine-readable definition files. Using IaC tools such as Terraform, Ansible, and CloudFormation, an organization can automate the setup and deployment of infrastructure environments, including high availability and disaster recovery strategies.

The main benefits of IaC for HA/DR include:

1) Consistency and Repeatability: Infrastructure deployment is consistent across all environments (development, staging, production).
2) Version Control: Infrastructure configurations are stored in version-controlled systems for rollback and auditing.
3) Scalability and Flexibility: Infrastructure can be modified and scaled based on changing needs quickly, with no manual intervention.

4) Faster Recovery: In the event of a disaster, IaC can be used to quickly recreate the required infrastructure and systems in a new region or availability zone resulting in significantly reduced recovery times.

VI. CLOUD SERVICE MODELS AND THEIR IMPACT ON HIGH-AVAILABILITY (HA) AND DISASTER RECOVERY (DR)

A. Infrastructure as a Service (IaaS)

IaaS offers virtualized computing resources, such as virtual machines (VMs), storage, and provisioning for networking, and enables users to design their own high availability/disaster recovery (HA/DR) configurations using the IaaS resources. HA/DR in IaaS typically involves:

1) Multiple Availability Zones (AZs): In platforms like AWS and Azure, you can deploy VMs across multiple AZs within a region to ensure availability even in the case of a failure in one zone.
2) Backup and Replication: Routine backups for Virtual Machines (VMs) and storage must be retained in remote locations or available in secondary regions. Replication of VMs and data can be achieved either synchronously or asynchronously, contingent upon the application’s requirements, service level agreements (SLAs), and recovery point objectives (RPO).

B. Platform as a Service (PaaS)

PaaS serves as an environment in which customers can build, deploy, and manage applications without concern for the underlying infrastructure. [7] In fact, the platform abstracts most of the complexity of HA/DR, but it is still necessary to design for HA/DR at the application level.

HA/DR in PaaS environments typically involves:

1) Built-in Redundancy: Cloud platforms like Google App Engine or Azure App Service automate distributing your app over multiple instances and availability zones.

   If an instance fails, another instance will take over the traffic and resolve downtime with a very short delay.
2) Automatic Failover: This technology is implemented in PaaS providers and minimizes user intervention during system outages, such as automatic traffic routing to healthy instances.
3) Geo-Redundancy: Some PaaS providers offer georedundancy, which refers to multiple regional instances enabled by the provider, allowing applications to run in different geographic localities. This is advantageous because traffic can be set up to redirect to the backup region in the event of a regional disaster.
4) Automatic Load Balancing: Most PaaS offerings balance load between multiple application instances automatically, allowing for high availability and scalability.
5) Database Redundancy: PaaS platforms like Azure SQL Database or Google Cloud SQL provide automatic replication and failover capabilities to ensure database availability, even during regional outages.

C. Software as a Service (SaaS)

SaaS providers oversee all infrastructure, platforms, and application stacks to provide a fully managed service. With respect to HA/DR in SaaS, the customer’s primary focus is whether the provider can maintain the service if the provider goes offline.

Key considerations include:

1) Provider-Side Redundancy: Major SaaS providers have their own high-availability and disaster recovery solutions, including geographically distributed data centers, replication, and backup.
2) Service Level Agreements (SLAs): As part of an organization’s cloud strategy, understanding the provider’s SLA is an important consideration. Providers such as Salesforce or Microsoft 365 offer SLAs for uptime and data recovery; but review the actual contract for the details around recovery time and guarantees.
3) Data Backups: While the provider generally provides some level of data backup and data replication services, it is important for organizations to consider the granularity and frequency of backups to determine whether a

third-party backup service is needed for additional security.

VII. APPLICATION TIERS: PLATINUM, GOLD, SILVER, AND BRONZE

In the process of developing an availability and disaster recovery (DR) strategy, it is important to categorize applications based on their criticality to the business [8]. This classification determines how much time and resources should be devoted to enabling each application to meet its recovery objectives (RTO/RPO) and service level agreements (SLAs). Tiers, such as Platinum, Gold, Silver, and Bronze, will be developed to facilitate resource allocation and to develop appropriate HA/DR strategies for each tier.

A. Platinum Tier: Mission-Critical Applications with Zero Tolerance for Downtime

Characteristics:

1) Zero tolerance for downtime: These applications are essential for the operation of the business, and any downtime can lead to severe financial loss, reputational damage, or regulatory compliance issues.
2) These systems must be available 24 hours per day, 7 days per week, 365 days per year and are expected to be restored immediately and with the least possible disruption to service.

Example Applications:

1) Financial systems (e.g., banking transactions, trading platforms)
2) Customer-facing platforms (e.g., e-commerce websites, payment gateways)
3) Healthcare applications (e.g., electronic health records, patient management systems)
4) Telecommunication systems (e.g., network monitoring systems)

HA/DR Considerations:

1) Multi-region and multi-AZ deployments for redundancy.
2) Real-time replication (synchronous) of data to minimize data loss.
3) Automated failover with minimal recovery time to ensure rapid continuity of service.
4) Continuous monitoring with proactive alerts and quick incident response protocols.
5) Zero-touch automation for failover and recovery processes to ensure they can happen without human intervention.

B. Gold Tier: High-Priority Applications with Minimal Acceptable Downtime

Characteristics:

1) High priority, but not as essential as Platinum: These applications are used in business operations, but brief downtimes (a couple of hours) are acceptable and will not result in disastrous outcomes.

Example Applications:

1) Sales platforms (e.g., CRM systems, point-of-sale systems)
2) Internal business tools (e.g., procurement systems, inventory management)
3) Marketing and analytics platforms (e.g., ad management tools, reporting dashboards) HA/DR Considerations:

1) Redundant systems across multiple availability zones to ensure high availability.
2) Asynchronous replication of databases, where data loss is acceptable but must be minimized.
3) Fast failover capabilities to ensure minimal downtime, typically aiming for recovery within 1-4 hours.
4) Backup systems that can restore applications quickly with pre-configured recovery points.
5) Manual failover processes could be part of this tier, depending on the system.

C. Silver Tier: Important but Less Critical Applications with Moderate Downtime Tolerance

Characteristics:

1) Moderate downtime tolerance: These applications are important but do not relate to direct business-critical functions. Hours or, in rare cases, one day of downtime is acceptable.

Example Applications:

1) HR systems (e.g., payroll, benefits management)
2) Internal support systems (e.g., IT help desk)

3) Knowledge management systems (e.g., document repositories)
4) Non-essential services that support internal functions but don’t have an immediate impact on revenue generation.

HA/DR Considerations:

1) Local redundancy with daily backups for data protection.
2) Failover may be manual or only needed during significant disruptions.
3) Moderate recovery objectives with a 4–12-hour RTO and longer RPO.
4) Data replication may be asynchronous, with acceptable delays in data synchronization.
5) Scheduled maintenance windows may allow for periodic downtime with minimal user impact.

D. Bronze Tier: Non-Essential Applications Where Downtime Is Acceptable and Cost Efficiency Is Prioritized

Characteristics:

1) Cost efficiency is the primary goal: These are nonessential, low-priority applications that assist with business functions but are not critical to business operations. Users not having access to these applications is not a major issue, as they can be down longer without a major impact on the business.

Example Applications:

1) Internal communication tools (e.g., chat systems, forums)
2) Legacy systems (e.g., old software not in active use)
3) Testing and development environments that don’t need to be highly available.
4) Non-business-critical applications (e.g., employee entertainment or volunteer portals).

HA/DR Considerations:

1) Minimal redundancy: These systems may only have basic backup solutions or be hosted on cost-effective single-instance servers.
2) Longer RTO and longer RPO, with recovery occurring over a 24+ hour period.
3) Manual restoration or recovery from backup, with no strict uptime requirements.
4) No real-time replication; backups may occur periodically (e.g., weekly or monthly) depending on the application’s importance.
5) Cost-efficient solutions (e.g., using cheaper storage, or not using auto-scaling or redundancy features).

VIII. SUMMERY OF APPLICATION TIERS

Tier	Criticality	Acceptable Downtime	HA/DR Focus
Platinum	Mission-critical applications	Zero downtime	Real-time replication, multi-region failover, zero-touch automation *Example: Financial systems*
Gold	High-priority applications	Minimal downtime	Redundancy, fast failover, asynchronous replication *Example: Sales platforms*
Silver	Important but less critical applications	Moderate downtime	Local redundancy, manual failover, backup systems *Example:HR systems*
Bronze	Non-essential applications	Acceptable downtime	Basic backups, cost-efficient solutions *Example: Internal communication tools*

TABLE I

HIGH AVAILABILITY AND DISASTER RECOVERY (HA/DR) TIER

CLASSIFICATION

IX. KEY CONSIDERATIONS IN IMPLEMENTING HA/DR IN CLOUD ENVIRONMENTS

A. Cost and Budgeting

When planning for High Availability (HA) and Disaster Recovery (DR) in cloud environments, balancing cost-efficiency with the required level of availability or redundancy is a critical consideration.

1) Balancing Between Cost-Efficiency and Availability Cloud providers offer pricing models with different service-level choices and availability guarantees. Although it may be tempting to choose the least expensive option, it is important to ensure that the necessary level of redundancy (e.g., multiple copies of data or failover across multiple regions) is maintained. [9] Multi-region deployment, for example, enhances reliability but increases costs.

2) Considerations for Cost Optimization

a) Failover Costs: Failover within a single region vs. across multiple regions (cross-region replication) can be a major cost driver.

b) Resource Scaling: Automatic scaling ensures availability during traffic spikes but can significantly increase costs if not properly managed.

c) Storage Redundancy: High-availability storage (multi-AZ or geo-replicated) is more expensive than standard storage.

3) Understanding Pricing Models for Cloud Services Each cloud provider offers a varied pricing structure, and understanding these is crucial for cost optimization.

a) Data Transfer Costs: Transferring data across regions or availability zones can incur charges. Efficient data flow planning can reduce unnecessary expenses.

b) Storage Costs: Maintaining backup data in multiple regions or using continuous replication may lead to high costs. Organizations must determine acceptable cost levels and define essential requirements.

B. Security and Compliance

An effective HA/DR strategy must conform to security best practices and adhere to regulatory and industry standards for securing sensitive data.

1) Identity and Access Management (IAM) Strict IAM policies should be in place to limit access to HA/DR infrastructure, applications, and systems, ensuring that only authorized personnel can manage resources.

2) Network Security Secure network environments should be established using Virtual Private Networks (VPNs), firewalls, and security groups to protect both primary and secondary environments.

3) Data Encryption Data in backup and failover systems should be encrypted both at rest (stored) and in transit (moving between locations). Best-practice encryption protocols should be followed.

4) Access Control Role-Based Access Control (RBAC) should be implemented to restrict access to sensitive backup data and prevent unauthorized changes to HA/DR configurations.

5) Compliance Organizations must adhere to industryspecific regulatory requirements such as GDPR, HIPAA, and SOC 2. While cloud providers offer compliance support, organizations must verify and ensure full compliance with applicable regulations.

C. Monitoring and Alerts

Regular monitoring of the HA/DR environment is essential to ensure that it functions as intended and that failures or potential problems are detected and resolved promptly.

1) Proactive Monitoring HA/DR systems should be monitored continuously, and event-driven alerts should be configured to identify risks before they escalate. Health checks and backup verifications should be scheduled regularly.

2) Key Monitoring Metrics and Alerts

a) Backup Success: Alerts should be configured to notify when backups fail or are incomplete.

b) Health Check Failures: System health checks

(e.g., latency, availability) should trigger alerts for anomalies.

c) Resource Utilization: Alerts should be set up to detect excessive resource consumption, which may indicate scaling issues.

3) Utilization of Monitoring Tools Cloud monitoring tools such as AWS CloudWatch and Azure Monitor should be leveraged to track system health and performance, ensuring timely issue detection and resolution.

X. CONCLUSION

As cloud computing continues to play an increasing role in business operations, ensuring high availability and effective disaster recovery strategies is essential to maintaining continuity and minimizing disruptions. [10] While the cloud provides advantages such as cost efficiency and scalability, ineffective management of high availability and disaster recovery solutions can be detrimental to business continuity.

1) Challenges in HA/DR Implementation Despite the benefits of cloud-based HA/DR solutions, organizations face several challenges, including:

a) Managing network latency and ensuring seamless failover across geographically dispersed cloud regions.

b) Maintaining data consistency across multiple regions while minimizing data loss.

c) Balancing the high costs associated with redundant infrastructure against business requirements.

d) Guaranteeing service levels and business continuity during extended outages.

2) Strategic Solutions for Businesses Businesses must develop strategic solutions that balance performance, availability, and recovery objectives. These solutions should address the challenges of HA/DR through optimized resource management and proactive planning.

3) Future Trends and Implications As cloud computing continues to evolve, organizations will need to adopt advanced technologies and operational approaches, including:

a) AI-driven disaster recovery solutions that predict failures and automate recovery actions.

b) Automation in failover systems to reduce human intervention and speed up recovery processes.

c) Hybrid or multi-cloud strategies to distribute workloads across multiple providers for better redundancy.

4) Evolving Industry Standards With the growing demand for high availability and disaster recovery, new industry standards and best practices are likely to emerge to support businesses in achieving resilience.

5) Conclusion and Business Imperatives To maintain sustainability and competitiveness in an ever-changing digital landscape, organizations must continuously innovate and enhance their cloud strategies. Ensuring security, resiliency, and adaptability in HA/DR solutions will be crucial for long-term success.

REFERENCES

[1] T. Wood, E. Cecchet, K. K. Ramakrishnan, P. Shenoy, J. van der Merwe, and A. Venkataramani, “Disaster recovery as a cloud service: economic benefits & deployment challenges,” in *Proceedings of the 2nd USENIX Conference on Hot Topics in Cloud Computing*, HotCloud’10, (USA), p. 8, USENIX Association, 2010.

[2] Y. Gu, D. Wang, and C. Liu, “Dr-cloud: Multi-cloud based disaster recovery service,” *Tsinghua Science and Technology*, vol. 19, no. 1, pp. 13–23, 2014.

[3] B. Liu, Y. Xin, and C. Zhang, “A solution for a disaster recovery service system in multi-cloud environment,” pp. 1–4, 12 2022.

[4] A. Abualkishik, A. Alwan, and Y. Gulzar, “Disaster recovery in cloud computing systems: An overview,” *International Journal of Advanced Computer Science and Applications*, vol. 11, p. 702, 10 2020.

[5] D. Clitherow, M. Brookbanks, N. Clayton, and G. Spear, “Combining high availability and disaster recovery solutions for critical it environments,” *IBM Systems Journal*, vol. 47, no. 4, pp. 563–575, 2008.

[6] C. Gabriel, *Data Center Disaster Recovery and High Availability*, pp. 639–657. 2015.

[7] J. Savill, *Backup, High Availability, Disaster Recovery, and Migration*, pp. 297–323. 2020.

[8] J. Mendonc¸a, E. Andrade, P. T. Endo, and R. Lima, “Disaster recovery solutions for it systems: A

systematic mapping study," *Journal of Systems and Software*, vol. 149, pp. 511–530, 2019.

[9] T. Adeshiyan, C. R. Attanasio, E. M. Farr, R. E. Harper, D. Pelleg, C. Schulz, L. F. Spainhower, P. Ta-Shma, and L. A. Tomek, "Using virtualization for high availability and disaster recovery," *IBM Journal of Research and Development*, vol. 53, no. 4, pp. 8:1–8:11, 2009.

[10] S. Tatineni, "Cloud-based business continuity and disaster recovery strategies," 11 2023.