International Journal of Scientific Research and Modern Technology (IJSRMT) ijsrmt.com
Volume 2, Issue 8, 2023
DOI: https://doi.org/10.38124/ijsrmt.v2i8.854

Architectural Approaches to Disaster Recovery and High Availability in SAP HANA Cloud

Gautam Yadav[1]

[1]Digital Transformation Leader, Kansas City, USA

Publication Date 2023/08/29

Abstract

The migration of mission-critical enterprise workloads to SAP HANA Cloud necessitates architectures that guarantee business continuity amidst disruptions. Downtime, whether caused by infrastructure failure, cyberattacks, or regional disasters, poses significant operational and financial risks. This paper provides a comprehensive analysis of architectural paradigms for disaster recovery (DR) and high availability (HA) in SAP HANA Cloud. We evaluate the trade-offs inherent in active-active, active-passive, and hybrid cloud DR models, considering critical metrics such as Recovery Time Objective (RTO), Recovery Point Objective (RPO), cost, and complexity. Synthesizing these findings, we propose an integrated reference architecture that harmonizes SAP HANA's native replication features with advanced hyperscaler resilience services. A real-world case study implementation on Microsoft Azure demonstrates the practical application of these principles, resulting in an RTO of <10 minutes and an RPO of near-zero while optimizing costs. Our findings indicate that a strategic, automated, and cloud-native approach is essential for building resilient, compliant, and cost-effective SAP HANA Cloud environments, and we offer recommendations for future evolution involving AI-driven orchestration and multi-cloud strategies.

Keywords: *SAP HANA Cloud, Disaster Recovery, High Availability, Cloud Architecture, Business Continuity, Hybrid Cloud, System Resilience, Database Replication, Enterprise IT, Digital Transformation.*

I. INTRODUCTION

Enterprise applications are the operational core of modern organizations, enabling real-time decision-making, continuous operations, and effective customer engagement. Among these, SAP HANA Cloud has become a leading in-memory, cloud-native platform, powering enterprise resource planning (ERP), analytics, and mission-critical workloads. Its ability to process massive datasets in real-time has established it as a cornerstone of digital transformation across industries, including manufacturing, finance, healthcare, and retail.

However, as businesses migrate critical processes to the cloud, the need for robust disaster recovery (DR) and high availability (HA) becomes paramount. Relying on traditional backup-and-restore strategies is no longer sufficient. In today's interconnected digital economy, downtime—caused by hardware failures, natural disasters, cyberattacks, or cloud outages—can result in substantial financial losses, reputational damage, customer churn, and regulatory non-compliance. With research indicating the average cost of IT downtime can range from thousands to millions of dollars per hour, proactive resilience planning is essential.

Ensuring resilience in SAP HANA Cloud environments requires a comprehensive architectural strategy. This involves not only data protection but also combining redundancy, system replication, automated failover, and proactive monitoring. Modern hyperscalers like Microsoft Azure, Amazon Web Services (AWS), and Google Cloud Platform (GCP) offer advanced tools—such as geo-redundant storage, availability zones, and cloud-native DR orchestration services—to build resilient, end-to-end solutions. The challenge for organizations is to evaluate these options and select an approach that optimally balances scalability, performance, compliance, and cost efficiency.

Furthermore, adherence to strict regulatory frameworks such as GDPR, HIPAA, and financial industry mandates is critical for maintaining data availability, geographic redundancy, and achieving required recovery objectives. Therefore, DR and HA solutions must be designed to meet both technical resilience and stringent compliance standards and service-level agreements (SLAs).

This paper explores various architectural approaches to disaster recovery and high availability in SAP HANA Cloud, analyzing models such as active-active clustering, active-passive failover, and hybrid cloud disaster recovery. A reference architecture is proposed to demonstrate best practices for achieving enterprise-grade resilience. A case study implementation on Microsoft Azure is then presented to highlight practical design choices and outcomes for IT leaders and architects. The paper concludes with recommendations for future-proofing enterprise SAP landscapes through automation, AI-driven monitoring, and hybrid multi-cloud strategies.

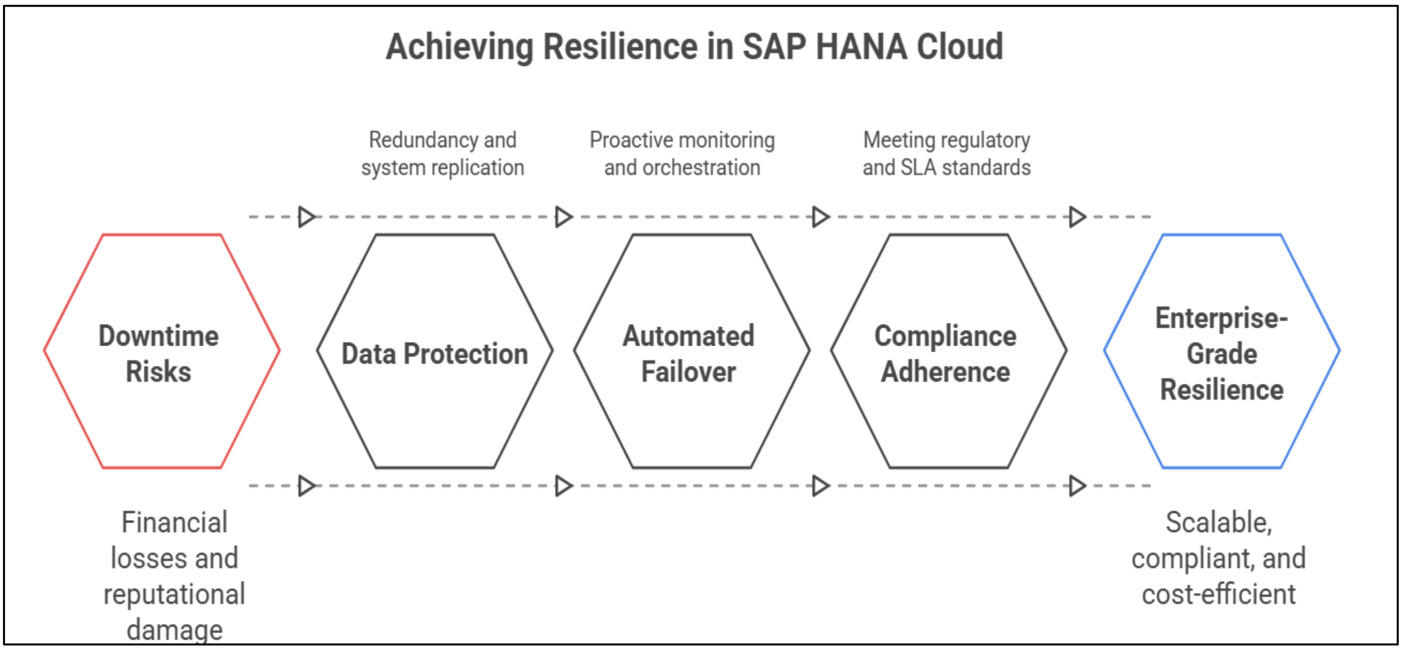

Fig 1 Achieving Resilience in SAP HANA Cloud

II. LITERATURE REVIEW

Resilience in enterprise IT, particularly concerning disaster recovery (DR) and high availability (HA), has been a pivotal area of research, evolving considerably with the transition of mission-critical workloads to cloud environments. This literature review traces the evolution of these strategies from traditional, on-premises paradigms to modern, cloud-native approaches, with a specific focus on their application and architectural implications within the SAP HANA ecosystem.

- *Evolution of Disaster Recovery and High Availability*

Historically, DR strategies for enterprise systems were resource-intensive and often complex. Research by Zhou et al. (2019) details the reliance on secondary data centers, manual tape backups, and asynchronous replication. While effective for data protection, these methods were associated with high costs, lengthy Recovery Time Objectives (RTOs), and potential for significant data loss, resulting in high Recovery Point Objectives (RPOs). High availability was primarily achieved through on-premises clustering solutions, necessitating substantial capital investment in specialized hardware and management expertise.

The advent of cloud computing fundamentally shifted the approach to resilience. Studies by Patel and Shah (2020) demonstrated how public cloud platforms, such as AWS, Azure, and Google Cloud, enable more agile and cost-effective recovery models. Key cloud features, including object storage for backups, availability zones for high availability, and automated provisioning, allowed organizations to move away from expensive standby data centers toward more flexible, on-demand, pay-as-you-go models. However, as noted by Liao et al. (2018), relying on static recovery plans is often insufficient for dynamic cloud environments, underscoring the need for more sophisticated, automated orchestration.

- *Resilience in On-Premises SAP HANA Environments*

Before the emergence of SAP HANA Cloud, research on SAP HANA resilience focused heavily on on-premises deployments. Müller et al. (2018) provided foundational work on achieving high availability using SAP HANA System Replication combined with operating system-level clustering tools like Pacemaker. These solutions ensured rapid failover within a single data center. However, deploying a robust, off-site DR solution for on-premises SAP HANA remained a complex undertaking, often requiring a second, fully-provisioned system and sophisticated management of data replication and application failover logic. The architectural complexity of integrating thousands of SAP application components with the HANA database layer presented a consistent challenge for comprehensive DR planning.

- *SAP HANA and Hyperscaler Integration for Resilience*

Recent literature has highlighted the combined power of SAP HANA's native features and hyperscaler cloud capabilities. Kumar and Singh (2021) showed how leveraging services like Azure Site Recovery or AWS Elastic Disaster Recovery alongside SAP HANA System Replication can dramatically reduce RTO and RPO for mission-critical workloads. In these hybrid-style architectures, enterprises can use cloud services to replicate the SAP application layer, while SAP HANA's built-in replication handles the database-level consistency. The SAP HANA Cloud architecture, a Platform-as-a-Service (PaaS) offering, abstracts much of this infrastructure complexity, providing built-in HA with automated failover and data replication across availability zones within a single region. However, a comprehensive DR plan extending beyond a single cloud region still requires a thoughtful architectural strategy.

➢ *Hybrid Cloud and Multi-Cloud DR Strategies*

The push for greater resilience and avoidance of vendor lock-in has led to increased interest in hybrid cloud and multi-cloud strategies. Wang and Chen (2022) explored hybrid approaches where primary workloads run in one cloud region or on-premises, with DR systems in a separate region, often leveraging cloud services for storage and compute. This ensures geographic redundancy and provides a balanced approach to control and cost. Multi-cloud strategies, involving two or more cloud providers, further distribute risk but introduce new complexities in managing data synchronization and recovery processes. Industry reports, such as those from Gartner (2023), emphasize that automation and intelligent recovery orchestration are critical for minimizing Mean Time to Recovery (MTTR) and navigating the complexities of these distributed landscapes.

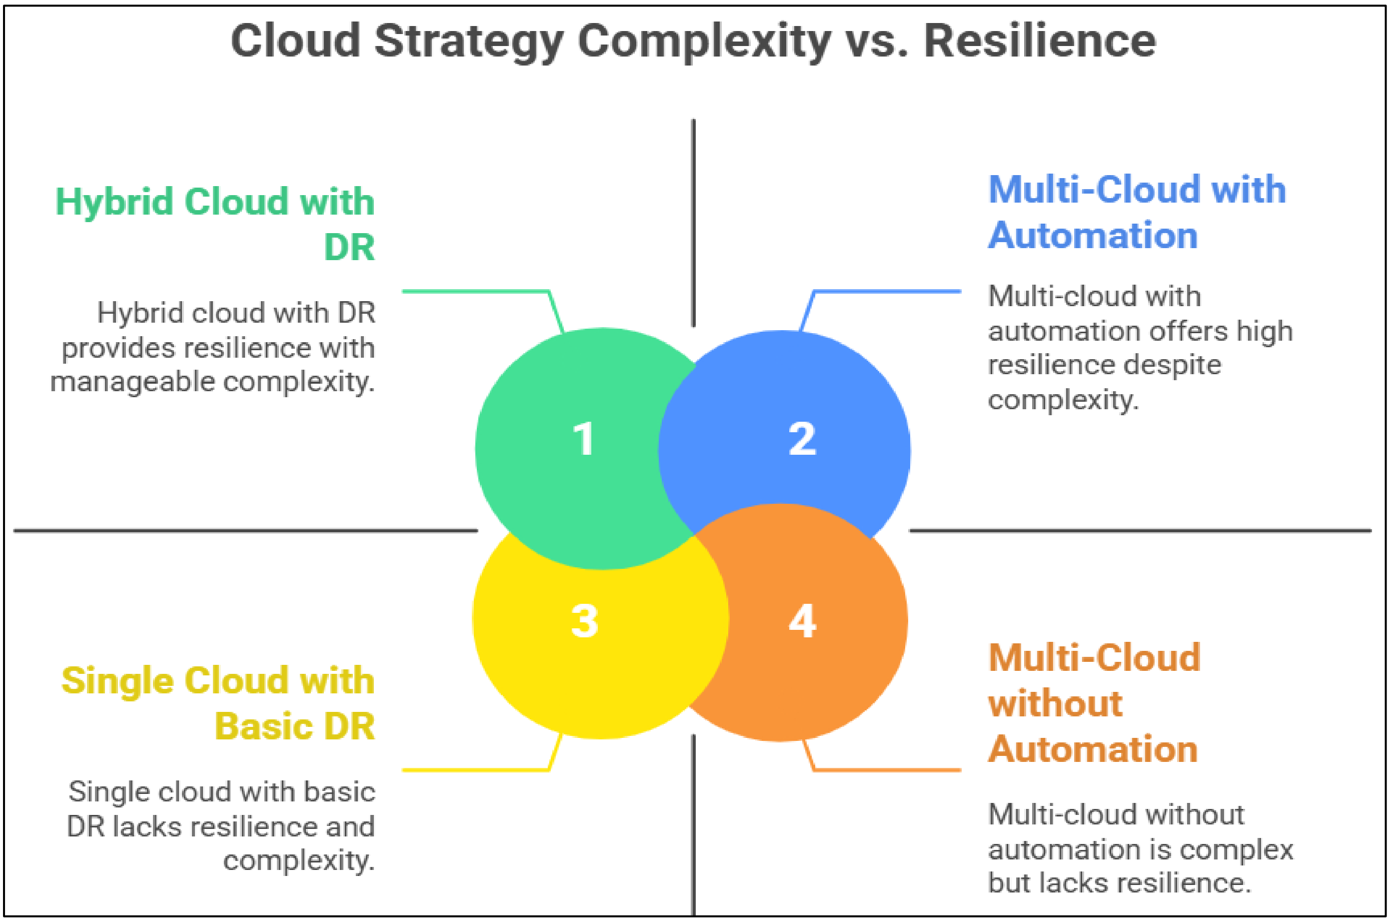

Fig 2 Cloud Strategy Complexity vs. Resilience

➢ *Research Gaps and Contribution of this Study*

Despite the significant body of work, several gaps remain, particularly concerning integrated architectural guidance for the fully managed SAP HANA Cloud environment.

- *Integrated Architectural Perspective:*

Existing studies often focus on either infrastructure-level resilience (e.g., hyperscaler features) or database-level replication (e.g., SAP HANA System Replication) in isolation. There is a lack of holistic, enterprise-grade reference architectures that integrate the native resilience of the SAP HANA Cloud platform with the advanced DR capabilities of hyperscalers.

- *Orchestration in Fully Managed Environments:*

The specific challenges and best practices for orchestrating failover and failback processes across the entire SAP application stack—especially in a fully managed PaaS context—have not been fully addressed.

- *Real-World Case Studies:*

While theoretical frameworks exist, there is a limited number of published case studies that provide empirical validation and actionable insights derived from real-world deployments of SAP HANA Cloud DR and HA solutions.

This paper addresses these gaps by providing an integrated architectural perspective. It synthesizes the evolution of DR/HA, incorporates recent advancements in SAP HANA Cloud and hyperscaler capabilities, and presents a comprehensive reference architecture. The practical case study on Microsoft Azure provides valuable deployment insights, bridging academic theory with practical application for IT leaders and architects navigating the complexities of SAP HANA Cloud resilience.

III. PROPOSED ARCHITECTURAL APPROACHES

➢ *Active-Active Architecture*

In a true active-active configuration for enterprise applications, two or more fully functional and independent SAP HANA Cloud landscapes are deployed and configured to operate simultaneously, sharing the workload. Both sites remain online, processing read-only transactions and potentially read-write queries in parallel, with data replicated and synchronized in real-time. For a database-centric platform like SAP HANA Cloud, this is most commonly implemented as an active-active (read-enabled) system replication, where a read-only replica can offload analytics workloads. Client connections are managed by a load balancer, which routes traffic to the optimal instance based on predefined rules. In the event of a failure, the load balancer automatically directs all traffic to the remaining healthy node(s), providing immediate, uninterrupted service.

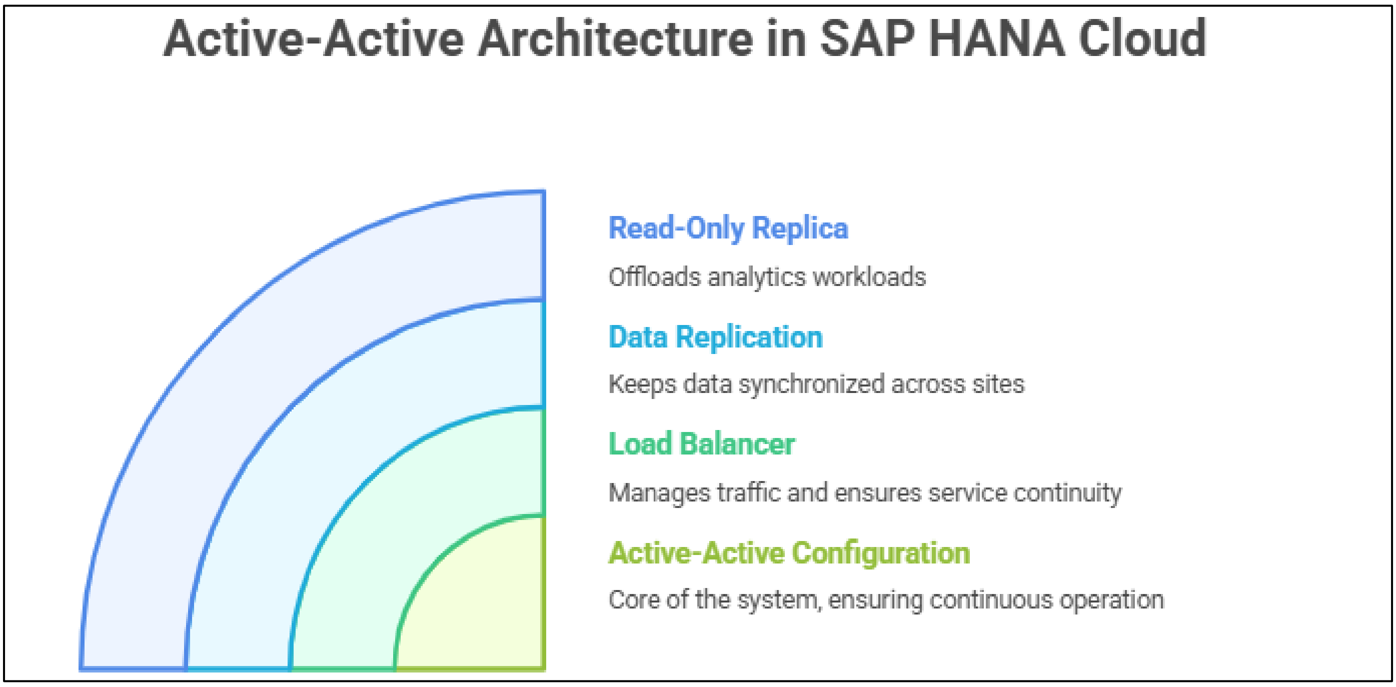

Fig 3 Active-Active Architecture in SAP HANA Cloud

- *Advantages:*

✓ *Continuous Availability:*

This model provides the highest level of availability, approaching a near-zero RTO. Failover is instantaneous and transparent to end-users, ensuring minimal disruption to business operations.

✓ *Enhanced Performance and Resource Utilization:*

By actively using all instances, this architecture distributes the workload, leading to improved performance, especially for read-intensive tasks. It also maximizes the utilization of expensive computing resources, offering a better return on investment compared to idle standby systems.

✓ *Zero Data Loss:*

Synchronous replication between sites ensures a near-zero RPO. Transactions are committed on both the primary and secondary instances before the application receives a confirmation, guaranteeing that no data is lost during a failover event.

✓ *Improved Scalability:*

The parallel processing capability allows for greater scalability, enabling the system to handle increasing transaction volumes and user load without performance degradation.

- *Challenges:*

✓ *Higher Cost:*

This is the most expensive architectural model due to the requirement for a complete, duplicate set of infrastructure, licenses, and potentially higher networking costs for synchronous replication.

✓ *Operational Complexity:*

Managing and monitoring an active-active environment is inherently more complex. It requires sophisticated synchronization mechanisms, robust monitoring tools, and careful management of potential data consistency conflicts, especially with read-write capabilities.

✓ *Application Constraints:*

Not all enterprise applications are designed to function seamlessly in an active-active setup, particularly those with complex transaction dependencies. The application layer must be stateless or manage sessions intelligently across nodes, and careful planning is required to avoid data conflicts from simultaneous write operations.

✓ *Network Dependency:*

The reliability and performance of synchronous replication are highly dependent on a high-speed, low-latency network connection between the two sites, which can be a limiting factor, especially across geographical regions.

- *Suitability:*

This model is best suited for enterprises with mission-critical SAP applications that have an extremely low tolerance for downtime. Industries such as banking, finance, healthcare, and telecommunications, where continuous, 24/7 access to services is a non-negotiable requirement, are ideal candidates for an active-active architecture. It is a strategic choice for organizations where the financial and reputational costs of even a brief outage far outweigh the higher initial investment and operational complexity.

➢ *Active-Passive Architecture*

In an active-passive configuration for SAP HANA Cloud, one system operates as the primary, handling all production workloads. A second, identical system remains in a "hot standby" mode, continuously updated through asynchronous or synchronous replication. The standby system is idle during normal operations but is fully prepared to take over if the primary system fails. A failover mechanism, which can be automated or manual, is responsible for detecting the primary system's failure and promoting the secondary to the active role. This model is a popular choice for balancing resilience with cost-efficiency.

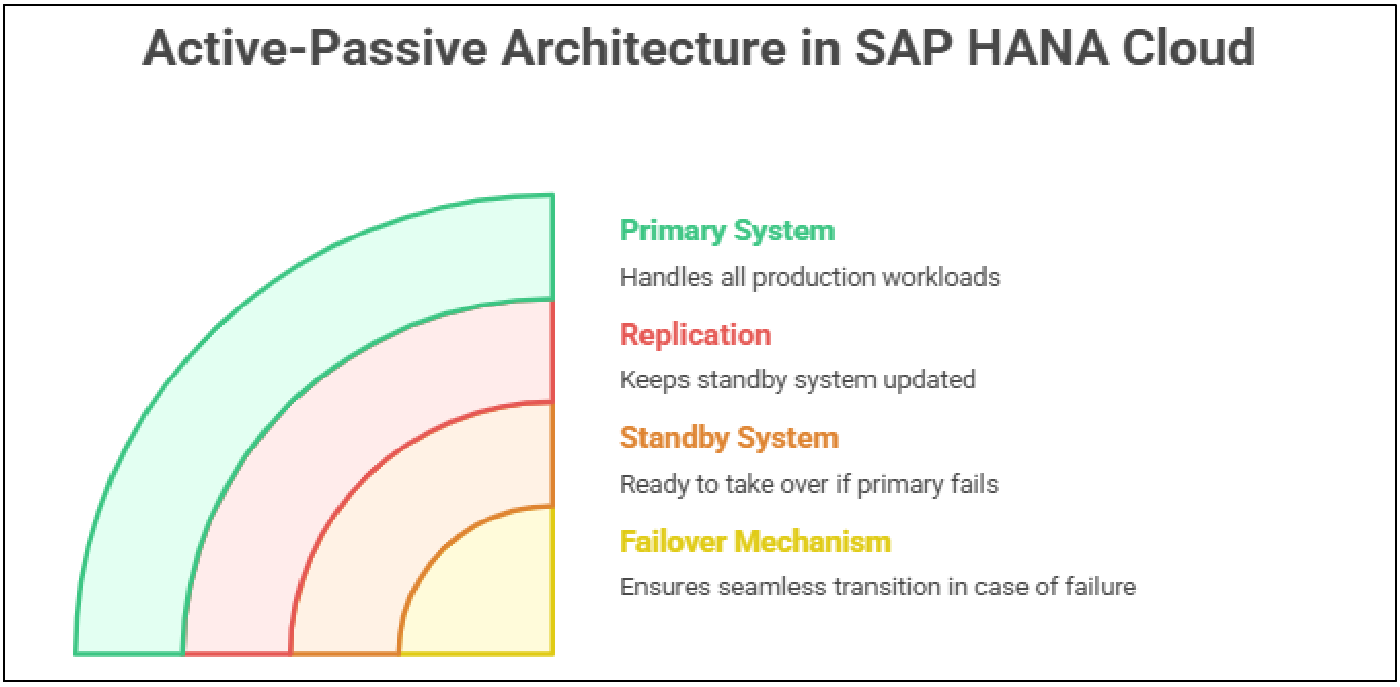

Fig 4 Active- Passive Architecture in SAP HANA Cloud

- *Advantages:*

✓ *Cost-Effective:*

This is a more budget-friendly option than an active-active setup, as the passive system is typically not used for production workloads during normal operations. This can reduce licensing costs and ongoing resource consumption, as the standby instance may be smaller or run in a minimal state until failover.

✓ *Simpler Administration:*

The active-passive model is less complex to configure and manage. Since only one system is processing transactions at any given time, there are fewer issues with data consistency and synchronization, eliminating the need for complex conflict resolution mechanisms.

✓ *Predictable Failover:*

The failover process is well-defined and predictable. Organizations can conduct regular failover and failback drills to ensure operational readiness, making recovery procedures more reliable and easier to execute.

✓ *Suitability for Traditional Workloads:*

This architecture is well-suited for SAP applications that cannot easily be horizontally scaled or have a primary-replica dependency, where all writes must go to a single point.

- *Challenges:*

✓ *Brief Downtime During Failover:*

While the downtime is significantly lower than a traditional backup-and-restore process, there is a short, but measurable, interruption during the switchover from the primary to the passive system. This brief service interruption occurs while the failover mechanism detects the failure, promotes the secondary instance, and reroutes network traffic.

✓ *Resource Underutilization:*

The primary disadvantage of this model is the underutilization of resources. The passive system remains idle and does not contribute to the processing workload during normal operation, which can be an inefficient use of a substantial hardware or cloud investment.

✓ *Data Loss Potential (Asynchronous Replication):*

For disaster recovery purposes where the standby is in a different geographical region, asynchronous replication is often used. This method can result in some minimal data loss (measured by the RPO) if the primary system fails before the most recent transactions have been replicated to the standby.

✓ *Limited Scalability:*

Unlike an active-active model, this architecture does not natively provide linear scalability. The performance and capacity of the entire system are determined by the capabilities of the single active node, which could become a bottleneck under heavy load.

- *Suitability:*

The active-passive model is a robust choice for organizations that need a high level of resilience for their SAP HANA Cloud workloads but can tolerate a brief period of downtime during a failover event. It is particularly relevant for scenarios where budget constraints make a full active-active setup unfeasible. Industries that rely on mission-critical but non-hyper-scale applications often leverage this model to balance strong data integrity with a straightforward, cost-effective, and reliable failover process.

➢ *Hybrid Cloud Disaster Recovery*

The hybrid cloud disaster recovery model represents a strategic compromise between the traditional control of on-premises environments and the flexibility and cost-effectiveness of the public cloud. In this architecture, an organization runs its primary SAP HANA Cloud system in one location (either on-premises, a private cloud, or a primary public cloud region) while leveraging another public cloud region as the designated disaster recovery site. During normal operations, data from the primary SAP HANA Cloud instance is continuously replicated to the secondary cloud location using asynchronous replication, ensuring geographic redundancy. In the event of a disaster affecting the primary site, the secondary public cloud environment is activated, and services are restored using the replicated data. This "pilot light" or "warm standby" approach provides a robust and resilient solution.

Fig 5 Hybrid Cloud Disaster Recovery Process

- *Advantages:*

✓ *Significant Cost Savings:*

The hybrid model provides substantial cost efficiency by minimizing capital expenditure. Instead of maintaining a duplicate, fully-sized DR data center, organizations only pay for the cloud resources they consume during replication and during a disaster recovery event or a test drill. This pay-as-you-go model contrasts sharply with the high costs of building and maintaining a traditional secondary site.

✓ *Geographic Redundancy and Risk Mitigation:*

By replicating data to a geographically separate public cloud region, the hybrid model protects against region-wide outages, natural disasters, and other localized threats that could affect a single data center. This ensures that business operations can continue even if the primary site is completely compromised.

✓ *Flexibility and Scalability:*

This architecture offers exceptional flexibility, allowing enterprises to scale their recovery environment on-demand. Resources can be rapidly provisioned and scaled up in the cloud during a failover and then scaled back down after a failback, aligning resource usage with actual need. This is a critical advantage for handling unforeseen spikes in workload demand during a disaster.

✓ *Simplified Testing:*

Testing a disaster recovery plan is far simpler and less disruptive in a hybrid environment. A sandbox environment can be spun up in the public cloud to test recovery procedures without impacting the production system. This allows for frequent testing to validate the DR plan, a critical best practice often neglected in traditional setups.

✓ *Compliance Support:*

For industries with strict data sovereignty requirements, the hybrid model allows sensitive SAP data to be kept on-premises or in a private cloud, while less sensitive information or backups can be securely replicated to the public cloud.

- *Challenges:*

✓ *Network Performance Dependency:*

The performance of the DR solution is highly dependent on the network connection and bandwidth between the primary and secondary sites. High network latency can increase the RPO, meaning more data could be lost during a failover. For synchronous replication, which requires low latency, geographic distance can be a significant constraint.

✓ *Complex Failover and Failback Orchestration:*

A well-defined and automated orchestration plan is crucial for a successful hybrid cloud DR strategy. Failover involves more than just the database; it requires activating the SAP application layer, rerouting network traffic, and ensuring all dependent services are available in the cloud. Manual failovers increase the RTO and risk of human error.

✓ *Security and Compliance Gaps:*

Hybrid environments introduce additional security and compliance challenges, as data needs to be securely managed and encrypted both on-premises and in the cloud. Ensuring consistent identity and access management across both environments is a key concern.

✓ *Inconsistent Toolsets:*

Managing resources across different environments (on-premises and a public cloud) can be complex, often requiring different toolsets and skill sets for administration.

- *Suitability:*

The hybrid cloud DR approach is an excellent choice for organizations seeking to reduce the capital expense of a dedicated DR site while simultaneously enhancing their resilience posture. It is particularly valuable for enterprises managing sensitive or legacy on-premises SAP systems alongside their modern SAP HANA Cloud instances, allowing for a phased transition to a cloud-first DR strategy. This model ensures business continuity while providing the cost control and flexibility needed for evolving business requirements.

➢ *Comparison of Architectural Approaches*

In comparing architectural approaches for disaster recovery and high availability in SAP HANA Cloud, it's essential to consider the trade-offs across several critical dimensions: Recovery Point Objective (RPO), Recovery Time Objective (RTO), cost, complexity, and the best-suited use case.

Here is a revised comparison table that expands on the provided information, offering more detailed nuances for each category:

Approach Comparison

Approach	RPO	Cost	Complexity	Best Use Case
Active-Active	Near zero	Near zero	Very High	Critical Industries
Active-Passive	Low (minutes)	Low (minutes)	Moderate	Enterprises
Hybrid Cloud DR	Variable (minutes to hours)	Variable (minutes to hours)	Moderate-High	Organizations

Fig 6 Approach Comparison

➢ *Elaborated Comparison Points*

- *RPO (Recovery Point Objective):*

This metric defines the maximum acceptable amount of data loss following a disaster.

✓ *Active-Active:*

Achieves near-zero RPO by employing continuous, often synchronous, data replication across active sites. According to Google Cloud documentation synchronous replication can be used within a region for instances that reside in any zone.

✓ *Active-Passive:*

Typically achieves a low RPO. While synchronous replication can yield a near-zero RPO, asynchronous replication, often used for geographically distant standby sites, may result in a minimal loss of data (e.g., minutes of transactions).

✓ *Hybrid Cloud DR:*

The RPO is highly variable and depends on the chosen replication strategy (synchronous vs. asynchronous) and the distance between the primary and DR sites. Asynchronous replication is common for cross-region disaster recovery, which can result in an RPO of minutes or potentially hours if replication lags significantly.

- *RTO (Recovery Time Objective):*

This metric defines the maximum tolerable downtime of an application after a disaster.

✓ *Active-Active:*

Aims for near-zero RTO, as both systems are live, and failover is typically instantaneous, handled by workload routing mechanisms like load balancers.

✓ *Active-Passive:*

Achieves a low RTO, generally in the range of minutes, depending on the automated failover process, the time needed for the passive system to be promoted, and network redirection.

✓ *Hybrid Cloud DR:*

The RTO can range from minutes to several hours, influenced by factors like the automation level of the recovery process, the time needed to provision cloud resources, and the complexity of bringing the entire application stack back online.

- *Cost:*

This refers to the financial investment required for infrastructure, licenses, data transfer, and operational overhead.

✓ *Active-Active:*

Incurs the highest cost due to the need for fully redundant and operational infrastructure across multiple sites.

✓ *Active-Passive:*

Presents a medium cost profile, as the standby infrastructure might be smaller or used for non-production purposes, but it still requires dedicated resources and replication costs.

✓ *Hybrid Cloud DR:*

Offers the potential for the lowest cost, as DR resources in the cloud can be minimized ("pilot light" or "warm standby") and spun up only when needed. However, network egress costs and the complexity of integration can influence the overall cost.

- *Complexity*

This assesses the difficulty involved in designing, implementing, managing, and maintaining the DR/HA solution.

✓ *Active-Active*

Highest complexity, requiring sophisticated data synchronization mechanisms, conflict resolution, global load balancing, and stringent monitoring to ensure consistent state across active instances.

✓ *Active-Passive:*

Moderate complexity. While simpler than active-active, it still requires robust replication, a reliable failover mechanism, and well-defined testing procedures.

✓ *Hybrid Cloud DR:*

Can range from moderate to high complexity, depending on the level of integration between on-premises and cloud environments, the tools used for replication and orchestration, and the need to manage different security and networking models.

- *Best Use Case:*

Indicates the scenarios where each approach provides the most optimal balance of the above factors.

✓ *Active-Active:*

Ideal for applications where any downtime is unacceptable and where the highest level of performance and data consistency is required, such as real-time financial trading systems or critical patient care applications. Research Gate highlights that this approach maximizes load distribution and efficient resource utilization.

✓ *Active-Passive:*

Suited for organizations seeking a strong level of resilience for their critical SAP applications without the high cost of an active-active setup, accepting minimal downtime during a failover event.

✓ *Hybrid Cloud DR:*

Best for enterprises transitioning to a cloud-first DR strategy, those with geographically dispersed operations, or those needing to balance cost savings with compliance and resilience goals.

This comparison provides a framework for evaluating the most appropriate architectural approach for SAP HANA Cloud based on specific business requirements and risk tolerance. Trilio emphasizes that the key is to strike a balance between preventing disruptions and preparing for potential disasters. The optimal choice will depend on a detailed assessment of each organization's unique needs.

IV. REFERENCE ARCHITECTURE

To build a resilient SAP HANA Cloud architecture, seamlessly integrate its native replication with hyperscaler services and intelligent orchestration. A robust design hinges on four pillars: A Primary Site for production, a dedicated Secondary Site for disaster recovery, an optional Hybrid Extension for multi-cloud flexibility, and a centralized Monitoring & Orchestration layer.

The core of this resilience is SAP HANA System Replication (HSR), which synchronizes data in near real-time to minimize loss and meet strict compliance goals. This is empowered by automated failover tools from hyperscalers and SAP, enabling rapid recovery and operational continuity with minimal manual effort.

➢ *Primary Site:*

The active production environment.

➢ *Secondary Site:*

A dedicated disaster recovery (DR) environment.

➢ *Hybrid Extension (Optional):*

For cross-region or multi-cloud deployments.

➢ *Monitoring & Orchestration:*

A centralized layer for management and automation.

SAP HANA System Replication (HSR) serves as the foundational data synchronization mechanism, ensuring near-zero data loss and supporting stringent RPOs. For rapid recovery, automated failover orchestration—leveraging hyperscaler and SAP-native tools—minimizes downtime (RTO) and ensures business continuity without significant manual intervention.

- *Core Technologies:*

✓ *SAP HANA System Replication (HSR):*

Provides near real-time data synchronization to minimize data loss (low RPO).

✓ *Automated Failover Tools:*

Hyperscaler and SAP services enable rapid recovery, ensuring business continuity with a low recovery time objective (RTO).

V. CASE STUDY: SAP HANA CLOUD DR ON MICROSOFT AZURE

To illustrate the architectural approaches, this case study presents a disaster recovery (DR) and high availability (HA) design for SAP HANA Cloud deployed on Microsoft Azure. Business Context A global manufacturing enterprise migrated its mission-critical SAP S/4HANA applications to SAP HANA Cloud on Azure to improve scalability and reduce infrastructure costs. Ensuring business continuity was a top priority due to strict SLAs and compliance requirements.

➢ *Implementation Results:*

- *Recovery Time Objective (RTO):*

Under 10 minutes.

- *Recovery Point Objective (RPO):*

Near zero.

- *Cost Optimization:*

Active-passive minimized expenses while ensuring resilience.

- *Resilience Achieved:*

Operations continued during a simulated outage with no user impact.

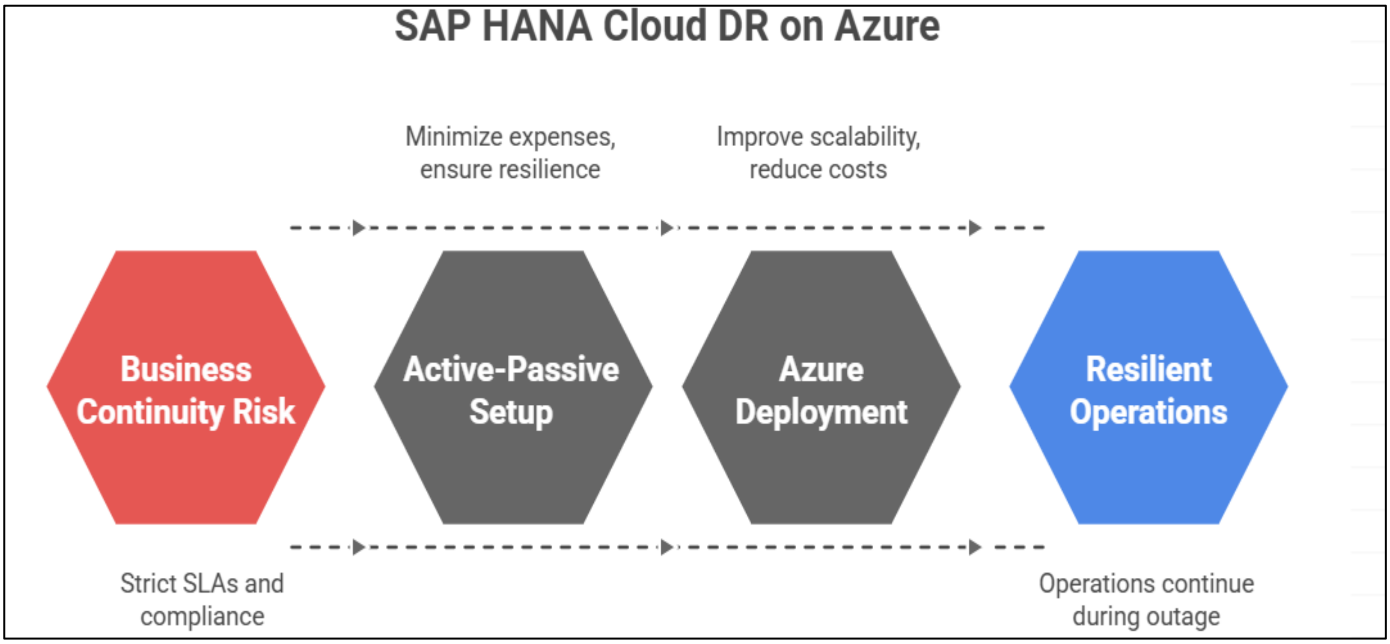

Fig 7 SAP HANA Cloud DR on Azure

VI. CONCLUSION AND FUTURE WORK

Ensuring disaster recovery (DR) and high availability (HA) is fundamental for enterprises operating business-critical workloads on SAP HANA Cloud. This paper examines architectural approaches, including active-active clustering, active-passive replication, and hybrid cloud disaster recovery, each offering unique benefits and trade-offs.

Key findings highlight that active-active models deliver maximum resilience but at high cost, active-passive offers a balanced approach, and hybrid cloud DR enables cost flexibility. Future work should explore automation with AI/ML, multi-cloud architectures, compliance-driven designs, and sustainability in DR/HA.

By adopting these approaches, enterprises can ensure that their SAP HANA Cloud environments remain resilient, compliant, and future-ready.

➢ *Acknowledgment*

The author would like to express sincere gratitude to colleagues, mentors, and the broader SAP community for their insights into disaster recovery and high availability practices. Special thanks to industry experts and cloud service providers whose reference architectures and best practices informed the development of this paper.

REFERENCES

[1]. Müller, T., Schmidt, A., & Hoffmann, R. (2018). System replication and high availability in SAP HANA environments. Journal of Enterprise Computing, 22(3), 145–159.

[2]. Zhou, Y., Li, H., & Zhang, P. (2019). Disaster recovery strategies in cloud computing: A systematic review. International Journal of Cloud Computing, 10(1), 1–15.

[3]. Patel, R., & Shah, M. (2020). Cloud-native disaster recovery: Leveraging hyperscaler capabilities. International Journal of Computer Applications, 176(15), 23–30.

[4]. Kumar, S., & Singh, R. (2021). Business continuity in SAP HANA Cloud: Replication and failover strategies. Proceedings of the International Conference on Cloud Technologies, 101–110.

[5]. Wang, J., & Chen, L. (2022). Hybrid cloud architectures for enterprise disaster recovery. International Journal of Innovative Research in Computer Engineering, 12(2), 56–65.

[6]. Gartner. (2023). Market Guide for Cloud Disaster Recovery Solutions. Gartner Research Report.

[7]. SAP SE. (2023). SAP HANA Cloud – High Availability and Disaster Recovery Guide. SAP Technical Documentation.

[8]. SAP SE. (2022). SAP HANA System Replication. SAP Help Portal. https://help.sap.com/docs/SAP_HANA_PLATFORM/6b94445c94ae495c83a19e6158039c4a/6be6d5e1b2ce4365922e43187c70172e.html

[9]. Microsoft Learn. (2023). High availability of SAP HANA on Azure VMs. Microsoft Documentation. https://learn.microsoft.com/en-us/azure/sap/workloads/sap-hana-high-availability

[10]. Amazon Web Services. (2023). High availability and disaster recovery for SAP HANA on AWS. AWS Documentation. https://docs.aws.amazon.com/sap/latest/sap-hana/hana-ops-ha-dr.html

[11]. Google Cloud. (2023). SAP HANA disaster recovery planning guide. Google Cloud Documentation. https://cloud.google.com/sap/docs/sap-hana-dr-planning-guide

[12]. SAP SE. (2022). SAP HANA Cloud — High Availability and Disaster Recovery. SAP Community Blog. https://blogs.sap.com/2022/10/25/sap-hana-cloud-high-availability-and-disaster-recovery

[13]. Disaster Recovery Journal. (2023). Mastering Disaster Recovery for SAP and HANA – Ensuring Business Continuity. DRJ Journal. https://drj.com/journal_main/mastering-disaster-recovery-for-sap-and-hana-ensuring-business-continuity-in-the-face-of-disruptions/

[14]. NetApp. (2023). SAP HANA Disaster Recovery Solution Comparison. NetApp Documentation. https://docs.netapp.com/us-en/netapp-solutions-sap/backup/hana-dr-anf-dr-solution-comparison.html

[15]. SAP Community. (2022). Best Practices for HA/DR for SAP Workloads on AWS. SAP Blog. https://blogs.sap.com/2022/04/13/best-practices-for-ha-dr-for-sap-workloads-on-aws/

[16]. Alibaba Cloud. (2020). High Availability and Disaster Recovery of SAP HANA. Alibaba Cloud Blog. https://www.alibabacloud.com/blog/high-availability-and-disaster-recovery-of-sap-hana_596525