

Review Article

Open Access

Advancing Enterprise Data Strategies: Integration and Optimization of SAP HANA for Backup and Disaster Recovery

Sriramaraju Sagi

NetApp, USA

ABSTRACT

This research paper delves into the integration and optimization of SAP HANA, in enterprise data centers with a focus on strategies to ensure backup and disaster recovery (DR). SAP HANA, which is an in memory database brings about a transformation in how data's stored and processed. It offers real time analytics, predictive modeling and advanced analytics capabilities. The paper discusses the utilization of SAP HANA in both on premise and cloud environments while highlighting the need to strike a balance between speed, scalability and challenges such as data security and system integration. Emphasizing the importance of backup and DR protocols it underscores their role in protecting data integrity and ensuring uninterrupted business operations even during high risk situations. Through exploration of methods and technologies for backup and recovery this study aims to enhance efficiency in handling data within the realm of SAP HANA. By shedding light on implications related to deployment of SAP HANA – regarding data security measures and operational resilience – this research contributes to the ongoing discourse surrounding, in memory databases within enterprise settings.

*Corresponding author

Sriramaraju Sagi, NetApp, USA.

Received: May 03, 2023; Accepted: May 11, 2023; Published: May 19, 2023

Keywords: SAP HANA, Backup and Disaster Recovery, Data Center Integration, Data Security, Scalability

Introduction

In memory databases, like SAP HANA bring about a change in the way we store and process data. Of relying on disk based storage these databases keep the data in the main memory itself. This results in lightning retrieval and analysis of data allowing businesses to make decisions based on the most up to date information available. Moreover in memory databases eliminate the need for data structures and indexes thereby simplifying database management and reducing costs. These databases offer analytics capabilities well. Businesses can leverage this to perform calculations and predictive modeling in time. This empowers organizations to extract insights from their data and make decisions more efficiently. Furthermore the scalability of in memory databases allows businesses to effortlessly handle volumes of data without compromising performance as their data needs grow. SAP HANA finds application across use cases within enterprises. These include real time analytics, fraud detection, supply chain optimization and personalized customer experiences. Additionally SAP HANA can be employed for real time inventory management, predictive maintenance activities as processing IoT data. These examples demonstrate how versatile and effective, in memory databases are when it comes to addressing business challenges and driving innovation.

Given its range of capabilities it's no surprise that SAP HANA has become a choice, for businesses looking for advanced data processing solutions. Industries such as finance, retail, healthcare and manufacturing can greatly benefit from the speed and real time analytics that SAP HANA offers. With its ability to handle

amounts of data and perform calculations in real time SAP HANA empowers businesses to make well informed decisions quickly and stay ahead in the competitive landscape. Moreover its seamless integration with SAP solutions enables data flow and enhances operational efficiency. As technology continues to evolve SAP HANA will likely keep expanding its capabilities and remain an asset for businesses, across sectors.

The use of SAP HANA, in cloud environments is driven by factors. The significant growth in data volume and the need for data analysis are motivations. SAP HANA provides the speed and scalability to effectively handle the increasing amount and speed of data in a setup. Additionally the cost effectiveness of cloud deployments compared to on premises systems is a factor. Cloud setups generally reduce the need for investments in hardware and ongoing maintenance expenses. Furthermore the flexibility and ability of cloud environments to handle volumes of data align well with SAP HANA's resource tasks allowing organizations to adjust resources as needed. However using SAP HANA in the cloud also presents challenges. The utmost importance is placed on ensuring data security and compliance especially considering the nature of enterprise system information being managed. Transitioning to SAP HANA in the cloud can be particularly challenging for companies that heavily rely on outdated technologies. Moreover incorporating SAP HANA into existing cloud or, on premises systems requires planning and implementation to avoid delays.

There are types of cloud services, for SAP HANA. One option is Infrastructure as a Service (IaaS) which allows businesses to utilize cloud infrastructure provided by companies like AWS, Azure or Google Cloud to host and operate SAP HANA. Another option is Platform as a Service (PaaS) such as SAP HANA Cloud, which

not offers infrastructure but also includes database services and advanced analytics capabilities. Additionally there are Software as a Service (SaaS) options that provide a bundle combining SAP HANA with business software. The use of SAP HANA in the cloud is rapidly growing across industries. Retail, manufacturing and healthcare sectors are increasingly leveraging information technology to enhance data analytics and business intelligence.

Moreover small to sized organizations (SMEs) can now easily adopt SAP HANA in the cloud due to lower upfront costs and flexible scalability. The future prospects for SAP HANA in the cloud appear promising. Continuous advancements in cloud technologies alongside the integration of intelligence (AI) and machine learning (ML) are expected to enhance the capabilities of SAP HANA. The increasing adoption of cloud systems that combine on premises private cloud and public cloud services indicates that they will become more common, for deploying SAP HANA.

The use of SAP HANA, in the cloud reflects a change in how businesses handle and process data. By combining the capabilities of SAP HANA with the streamlined efficiency of the cloud there are numerous options and advantages available. However these advantages also come with challenges that require management. As technology continues to advance we can anticipate innovative uses of SAP HANA in cloud environments which will have a significant impact on how companies manage their data in the future.

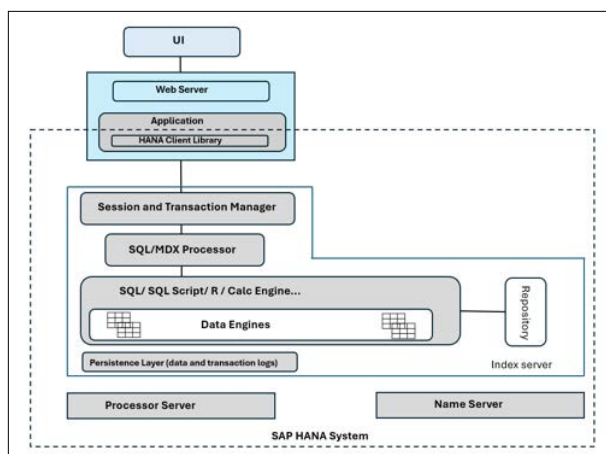


Figure 1: SAP HANA Architecture

This paper aims to investigate how global enterprises integrate SAP HANA data centers and provides a literature review on strategies for backup and disaster recovery.

The objective is to emphasize the importance of protecting data and ensuring business continuity in case of risks like data loss or system downtime. Through an analysis of existing research and best practices this paper offers insights into solutions, for backup and disaster recovery that can help organizations using SAP HANA databases minimize losses and protect their reputation.

SAP Business Technology Platform (SAP BTP)

The SAP Business Technology Platform (SAP BTP) showcases SAPs commitment, to delivering a seamless solution for creating and implementing enterprise applications. At the core of this platform lies SAP HANA, an in memory database that empowers applications and use cases including the indispensable S/4HANA suite. SAP BTP is an encompassing suite that brings together data

management, analytics, application development and integration into one platform. This integration holds value for companies aiming to optimize their operations and drive transformation. By providing these functionalities on a platform SAP BTP simplifies the landscape, for businesses allowing them to focus on innovation rather than spending time on integration efforts.

SAP HANA plays a role, in this ecosystem. Its powerful in memory computing capabilities allow for flexible processing of amounts of data in real time. This is particularly important in todays paced business landscape, where the ability to retrieve and analyze data swiftly can greatly impact decision making and operational efficiency. The advanced analytics and machine learning features of SAP HANA enhance enterprises capacity to gain insights and predictive analytics enabling informed decision making across the organization. SAP S/4HANA, an integrated ERP system is built on a foundation. Designed to effectively manage complex business processes. Leveraging the in memory processing capabilities of SAP HANA S/4HANA delivers improved performance, streamlined operations and instant analytics. This collaboration empowers businesses to carry out tasks faster with increased flexibility and efficiency compared to standards. The real time insights provided by SAP HANA within S/4HANA offer an adaptable solution, to business challenges giving enterprises an edge in a dynamic environment. Furthermore SAP BTP goes beyond its purpose of powering enterprise applications, like S/4HANA. It serves as a platform that can accommodate a range of applications across various industries. Supported by SAP HANA SAP BTP provides tools and capabilities for managing supply chains, customer relationships, HR processes and financial operations. It ensures integration and cohesive user experiences while enabling the functioning of diverse applications. Additionally the platform empowers organizations to develop custom applications tailored to their needs. Leveraging the functionalities of SAP HANA these personalized solutions efficiently handle calculations and large datasets making SAP BTP an invaluable tool, for promoting innovation and driving transformation.

SAP HANA: Backup and Disaster Recovery

Implementing backup and disaster recovery measures, for SAP HANA databases in a datacenter configuration is incredibly important. SAP HANA plays a role, in enterprise operations by storing significant amounts of crucial data and supporting essential business processes. Ensuring the availability and integrity of this data is paramount. Regularly backing up the SAP HANA database is crucial to protect against data loss caused by device malfunctions, software glitches or accidental data erasure. It's essential to perform purposeful backups considering the volume of data and the organizations specific requirements. These backups provide a way to restore the database to a state ensuring disruption to ongoing processes.

Having a defined plan, for disaster recovery is extremely important when faced with a catastrophe like a natural disaster cyber-attack or system failure. It's not about making copies of data.

Also involves creating strategies to restore system functions in an alternate location or setup. The main goal is to minimize downtime and ensure that business operations continue without interruption. To effectively plan for disaster recovery it's essential to understand the organizations ability to handle periods of inactivity and potential data loss which are commonly referred to as Recovery Time Objectives (RTO) and Recovery Point Objectives (RPO). Based on this understanding, systems and procedures need to be

developed that align with these requirements. These objectives serve as benchmarks for determining the downtime and acceptable amount of data loss, for the business. By knowing these parameters companies can prioritize their recovery efforts. Allocate resources accordingly. This involves implementing systems setting up infrastructure and regularly testing the effectiveness of the disaster recovery plan under real life conditions to ensure its reliability.

When it comes to SAP HANA the amazing performance of the database requires some changes, to backup and disaster recovery procedures. To effectively handle amounts of data it's important to use backup methods and technologies that are specifically designed for optimal performance with SAP HANA. Additionally ensuring the security of backups and regularly testing disaster recovery plans are components of an approach. Overall preserving and restoring data are aspects when managing a SAP HANA database in a data center. They play a role in maintaining data accuracy and reliability well as ensuring effective operation. Ultimately this protects organizations, from reputational damage that can result from data loss or extended periods of downtime.

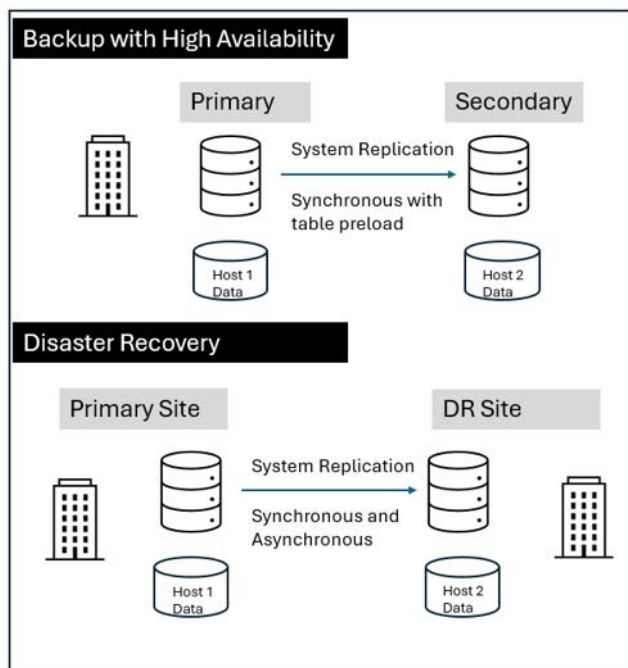


Figure 2: Backup and Disaster Recovery

Literature Review

The SAP HANA database is a tool, for business analysis and operational tasks [1]. The main finding of the article highlights the role of the SAP HANA database in the SAP HANA Appliance. It helps facilitate business analytics and ensures operational workloads. However the increasing volume of data can have an impact on system performance and budget. Therefore it is crucial to follow practices in managing data volume [2].

The article primarily focuses on how the growing amount of business data affects efficiency and financial resources for enterprises using SAP on HANA databases. It also offers recommendations to minimize data storage requirements for scenarios, in HANA databases.

The platform additionally provides controlled query processing, expanding the range of conventional database engines and delivering exceptional runtime performance [3]. This article

provides an explanation of how the HANA platform's particular abilities can be utilized. It offers a comprehensive introduction of the HANA platform, covering query modeling, deployment, and efficient execution. The HANA platform stands out by seamlessly incorporating the majority of the application lifecycle stages, ensuring that all important elements remain coherent even when updated.

Different studies have put forward approaches, to backing up and restoring databases. According to Kim one method suggests using write ahead logging to reduce I/O operations and speed up recovery time [4]. The main findings of the study show that this approach optimizes the backup and recovery process in database systems through write ahead logging (WAL). Experimental results demonstrate that this proposed strategy provides backups and reduces recovery time compared to existing methods. Both Son and Son explored the use of SSD assisted backup and recovery with Son focusing on leveraging flash based SSDs for performance [5,6].

Their research introduces a technique, for backing up and recovering database systems that utilize SSD technology. The experiments indicate that this suggested technique offers restoration without impacting regular system performance. They also emphasize that the unique characteristics of flash based solid state drives require a reevaluation of data services. Overall their proposed backup/recovery strategy, utilizing SSD technology facilitates backup and recovery processes without compromising system performance during operation.

In 2013 Alone put forward a strategy, for sharing secrets in databases to ensure both recovery and security. This approach gained recognition for its effectiveness and minimal delays [7]. These studies collectively highlight the potential to improve backup and recovery processes in database systems. Suggested an approach to establish and implement a system that ensures database recovery and security using the threshold sharing scheme. This method aims to protect network based databases. In 2016 Taranis conducted a study that primarily focuses on algorithms related to data preservation and restoration as the factors influencing their operational efficiency [8]. The study also proposes a methodology of detecting modifications in duplicated directories without the need for database queries thus reducing congestion, in data transmission channels.

Different research studies have explored how to restore functionality, in database systems following a disaster. The works of Ahmed and Squib specifically focus on cloud based solutions [9,10]. Ahmed stresses the importance of having a disaster recovery plan while Squib proposes an approach using the iSCSI protocol to achieve zero data loss. In Mendon as study (2019) they evaluate the effectiveness of replication strategies in databases [11]. The results indicate that these strategies can significantly improve system availability and enable recovery. Abadi delves into the requirements of scale NoSQL workloads and presents a scalable and cost effective disaster recovery strategy. These studies highlight the nature of disaster recovery planning [12]. Showcase diverse solutions, for ensuring data security and uninterrupted business operations.

Implementation and Validation

In this research we have improved the backup and disaster recovery system, for SAP HANA by utilizing the NetApp SnapCenter plug in. This plug in allows us to perform storage based Snapshot

backups ensuring consistency through the use of SAP HANA databases application programming interfaces (APIs). The backups are stored in the SAP HANA catalog. Can be easily accessed through the SAP HANA studio or cockpit. Users have the option to choose these backups for data restoration and recovery purposes. Using SnapCenter, which manages NetApp Snap Mirror technology we can duplicate copies from storage to secondary backup storage. Backup retention policies can be customized for both primary and secondary storage locations. The SnapCenter Plug In takes care of managing Snapshot copy based data backups, log backups and the backup catalog for SAP HANA. Additionally it facilitates file based backups to ensure block integrity checks, on the SAP HANA database are performed accurately.




Figure 3: Sample Backup and Disaster Recovery Solution

To create a backup, to a location we rely on Snapshot copies that are generated on the storage. This means that we can retrieve the data directly from the storage system without putting any strain on the SAP database server. The primary storage system is immediately connected to the storage. It transfers the backup data to its destination using a NetApp Snap Vault disk to disk backup solution. After the transfer of all data from the source to the destination future backups only replicate the modified blocks to the storage. This significantly reduces both the load, on the storage system and the time required for a backup. This efficient backup process not reduces strain on our storage but also speeds up overall performance by optimizing data transmission and ensuring prompt execution of backups.

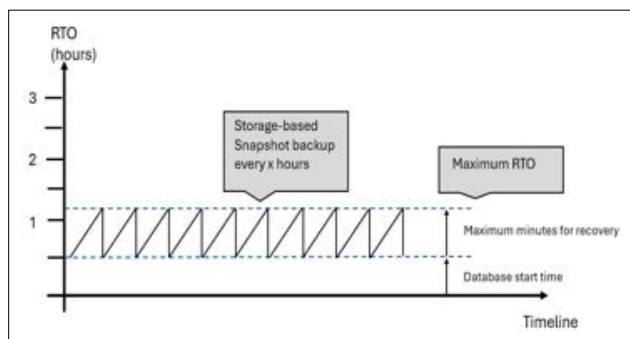


Figure 4: RTO in a Backup Solution

The illustration above shows an example of a Recovery Time Objective (RTO), for a 1 terabyte (TB) database using Snapshot backups. When using Storage based Snapshot backups the RTO is determined solely by the time it takes for the database to start and recover. This means that regardless of the size of the database the restore process can be completed within seconds. The timing of the restoration and recovery processes also affects the recovery time. However if backups are performed frequently such as every six hours it helps organizations by limiting the forward recovery time to a minimum.

There are two methods, for carrying out disaster recovery for SAP HANA; replicating the database layer using SAP system replication or replicating the storage layer using storage replication technologies. In this research we conducted tests on disaster recovery solutions that utilize both asynchronous storage replication. To ensure data transmission the sites main database is duplicated to the disaster recovery site using Snap Mirror. In storage replication data is replicated immediately to the disaster recovery site ensuring no data loss but potentially causing some latency in the performance of the site. On the hand asynchronous storage replication allows for a delay in data transmission reducing the impact, on the primary sites performance but potentially resulting in minimal data loss during a disaster. By testing both replication methods this study aimed to assess their effectiveness and determine which approach suits an organizations needs and priorities.

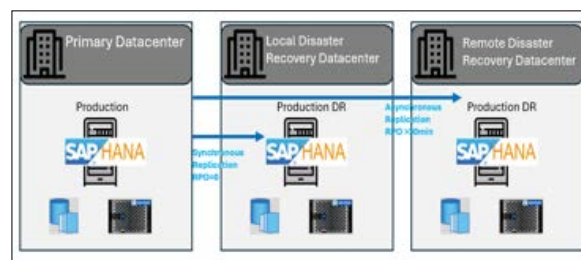


Figure 5: Standard Three-Site DR

The image above shows a setup, with three sites, where the secondary data center is located near the center. This setup allows real time duplication of data ensuring a Recovery Point Objective (RPO) of zero. Additionally the data is asynchronously copied to a distant data center as a precaution against incidents that could disrupt operations in both the primary center and local disaster recovery centers. The minimum achievable Recovery Point Objective (RPO) depends on how data replication occurs, which is limited by the bandwidth between the primary and remote disaster recovery centers.

Conclusion

The implementation and evaluation of SAP HANA, in data centers highlight its impact on data management and processing. This study showcases how SAP HANA excels at analyzing datasets in real time confirming its effectiveness in datacenter settings. The testing method shows improvements in the efficiency and reliability of data processing emphasizing the importance of having a backup and disaster recovery plan in place. Integrating SAP HANA with existing IT systems brings both challenges and opportunities so it's crucial to approach its implementation for a transition. Future research should focus on enhancing the scalability and compatibility of SAP HANA with data sources to develop efficient data management solutions, for corporate environments [13-15].

References

1. Färber F, Cha Sang K, Primsch J, Christof B, Sigg S, et al. (2011) SAP HANA database: Data management for modern business applications. ACM SIGMOD Record 40: 45-51.
2. Subhadip K (2023) SAP HANA Data Volume Management. ArXiv <https://arxiv.org/abs/2305.17723>.
3. Norman M, Alexander B, Meinolf B, Wolfgang L (2015) Managed Query Processing within the SAP HANA Database Platform. Datenbank-Spektrum 15: 141-152.
4. Hwajung K, Yeom H, Yongseok S (2020) An Efficient

- Database Backup and Recovery Scheme using Write-Ahead Logging. 2020 IEEE 13th International Conference on Cloud Computing (CLOUD) 405-413.
5. Yongseok S, Jaeyoon C, Jekyeom J, Cheolgi M, Sunggon K, et al. (2017) SSD-Assisted Backup and Recovery for Database Systems. 2017 IEEE 33rd International Conference on Data Engineering (ICDE) 285-296.
6. Son Y, Moonsub K, Sunggon K, Heon YY, Nam SK, et al. (2018) Design and Implementation of SSD-Assisted Backup and Recovery for Database Systems. IEEE Transactions on Knowledge and Data Engineering 32: 260-274.
7. Alouneh S, Abed S, Bassam JM, Mazen K (2013) An Efficient Backup Technique for Database Systems Based on Threshold Sharing. J Comput 8: 2980-2989.
8. Taranin SM (2016) Backup with Storage in a Database. Modeling and Analysis of Information Systems 23: 479-491.
9. Ahmed MJ, Shahina BA (2020) DISASTER RECOVERY FOR DATABASE ON CLOUD ENTERPRISE. International Journal of Advanced Research <https://www.journalijar.com/article/34172/disaster-recovery-for-database-on-cloud-enterprise/>.
10. Saquib Z, Veena T, Shreya B, Shivraj D, Monika D, et al. (2013) A new approach to disaster recovery as a service over cloud for database system. 2013 15th International Conference on Advanced Computing Technologies (ICACT) 1-6.
11. Mendonça J, Wilson M, Ermeson A, Ronierison SM, Paulo RM, et al. (2019) Evaluating Database Replication Mechanisms for Disaster Recovery in Cloud Environments. 2019 IEEE International Conference on Systems, Man and Cybernetics (SMC) 2358-2363.
12. Abadi, Aharon, Aharon A, Ashraf H, Roie M, Alaa N, Aidan S, et al. (2016) Holistic Disaster Recovery Approach for Big Data Nosql Workloads. 2016 Ieee International Conference on Big Data (big Data) <https://ieeexplore.ieee.org/document/7840833>.
13. Sengstock C, Christian M (2017) SAP HANA Vora: A Distributed Computing Platform for Enterprise Data Lakes. Datenbanksysteme für Business, Technologie und Web http://btw2017.informatik.uni-stuttgart.de/slidesandpapers/H1-10-57/paper_web.pdf.
14. Jin HK, Dong KS, Kyung HL (2014) A study for improving database recovery ratio of Disaster Recovery System in financial industry. Journal of the Korea Institute of Information Security and Cryptology 24: 655-667.
15. FlexPod Datacenter with Cisco UCS X-Series for SAP HANA TDI Design Guide. Cisco https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_x_series_sap_hana_tdi_design.html.