

# Large-Scale SAP HANA Database Backup and Restoration on AWS Cloud Infrastructure

Manoj Joshi

Dr. Babasaheb Ambedkar Marathwada University, India

---

## ARTICLE INFO

## ABSTRACT

Received: 05 Jan 2026

Revised: 08 Jan 2026

The migration of SAP HANA databases to AWS cloud infrastructure represents a fundamental transformation in enterprise resource planning systems management, enabling organizations to leverage elastic computing resources while maintaining mission-critical operational continuity. Cloud computing has evolved as a model for delivering computing resources as measured services over the internet, fundamentally changing how organizations approach infrastructure investments and operational scalability. SAP HANA's in-memory architecture presents unique challenges for backup and restoration operations on AWS, requiring specialized strategies that account for the database's columnar storage format, delta merge operations, and memory-optimized processing capabilities. AWS provides a comprehensive ecosystem of storage services, including Amazon S3 for object storage, Amazon EBS for block storage, and AWS Backup for centralized backup management, each offering distinct performance characteristics and cost structures suitable for different stages of the HANA backup lifecycle. The implementation of effective backup strategies for large-scale HANA databases on AWS demands careful consideration of backup windows, recovery time objectives, recovery point objectives, and regulatory compliance requirements while optimizing cloud resource consumption. This article examines architectural frameworks for implementing robust backup and restoration mechanisms for SAP HANA databases on AWS, analyzing the integration of HANA-native backup tools with AWS storage services, parallel processing capabilities through AWS infrastructure, and automation frameworks leveraging AWS management services. It addresses critical performance optimization techniques, including snapshot-based backups using Amazon EBS, incremental backup strategies utilizing HANA's delta backup capabilities, and lifecycle management policies that automatically transition backup data across Amazon S3 storage classes based on retention requirements and access patterns.

**Keywords:** SAP HANA, AWS Cloud, Database Backup, Amazon S3, EBS Snapshots, Disaster Recovery, Data Protection

---

## I. INTRODUCTION

Cloud computing represents a paradigm shift in information technology where computational resources are delivered as utility services over the internet, enabling organizations to provision and release resources dynamically based on demand patterns rather than maintaining fixed infrastructure capacity. The fundamental characteristic of cloud computing involves the transformation of computing infrastructure from capital expenditure models requiring substantial upfront investments in hardware and data center facilities to operational expenditure models where organizations consume resources on demand and pay only for actual usage [1]. This economic transformation has proven particularly significant for database systems where traditional deployments required organizations to provision infrastructure capacity for peak workload scenarios, resulting in substantial underutilization during normal operational periods. The elastic nature of cloud infrastructure enables database systems to scale computational and

storage resources dynamically in response to workload variations, optimizing resource utilization while maintaining performance characteristics required for enterprise applications [1].

SAP HANA represents a fundamental departure from traditional relational database management systems through its in-memory computing architecture, which maintains primary data structures in RAM rather than on disk storage systems. The in-memory architecture enables real-time analytics processing and transactional operations with response times substantially reduced compared to disk-based systems, transforming enterprise resource planning capabilities across industries. SAP HANA employs columnar storage formats optimized for analytical query patterns common in enterprise data warehousing scenarios, where queries typically access specific columns across large data sets rather than complete row records [2]. The columnar approach provides significant compression opportunities because data values within columns exhibit higher similarity compared to heterogeneous data types across rows, enabling compression algorithms to achieve substantial storage reduction ratios. The persistence layer in HANA maintains data durability through savepoint operations that periodically write in-memory data structures to disk storage, coordinating across all database services to ensure consistent snapshots despite continuous transaction processing [2].

Organizations implementing HANA on AWS infrastructure must address the complexity of protecting large-scale databases while maintaining production system availability, meeting stringent recovery time objectives, and optimizing cloud infrastructure costs through efficient resource utilization. AWS provides infrastructure specifically designed to support enterprise SAP workloads, with instance types optimized for HANA's memory-intensive operations and storage services architected to deliver performance characteristics required for large-scale database backup operations [3]. The AWS ecosystem for SAP HANA includes compute instances certified by SAP for production deployments, storage services spanning block and object storage architectures, networking capabilities enabling high-bandwidth data transfers, and management tools supporting automation of operational tasks including backup orchestration and disaster recovery procedures [3]. The integration capabilities between HANA's native backup utilities and AWS storage services enable organizations to implement sophisticated backup architectures that balance performance requirements during backup operations with cost optimization through intelligent data lifecycle management across storage tiers.

The architectural considerations for implementing HANA backup strategies on AWS extend beyond simple data protection to encompass broader operational requirements including high availability configurations using HANA System Replication across AWS Availability Zones, disaster recovery capabilities leveraging cross-region backup replication, and compliance frameworks that mandate specific retention periods and geographic data sovereignty requirements. AWS infrastructure provides programmatic control through APIs and command-line interfaces that enable automation of backup operations, integration with enterprise monitoring systems, and orchestration of complex recovery procedures that might involve coordinating database restoration with application server provisioning and network configuration. The challenge facing organizations lies in designing backup architectures that are native to AWS capabilities while respecting HANA-specific operational constraints including backup window limitations during peak transaction periods, memory resource allocation during backup operations, and coordination between data backups and log backup sequences that enable point-in-time recovery capabilities.

## **II. AWS INFRASTRUCTURE ARCHITECTURE FOR SAP HANA DATABASE SYSTEMS**

The computational infrastructure supporting SAP HANA on AWS encompasses specialized instance types designed to deliver the memory capacity, processing power, and I/O characteristics required for in-memory database operations. Cloud infrastructure exhibits considerable performance variability across different service providers and even within a single provider's offerings, with computational performance, network throughput, and storage I/O characteristics varying based on underlying hardware configurations, virtualization overhead, and resource contention from co-located workloads [4]. Performance studies of public cloud infrastructure have demonstrated that even instances with identical specifications can exhibit performance variations during different time periods, attributable to factors including network congestion patterns, storage backend load conditions, and thermal throttling of processors under sustained high utilization scenarios [4]. Organizations deploying HANA databases on

AWS must account for this performance variability when designing backup strategies, ensuring that backup window calculations incorporate appropriate margins to accommodate performance fluctuations that might extend backup completion times beyond typical execution durations.

The storage architecture for HANA on AWS employs Amazon EBS volumes configured to deliver the throughput and IOPS characteristics required for production database operations. EBS provides persistent block storage that operates independently of EC2 instance lifecycles, enabling data persistence across instance terminations and supporting snapshot capabilities for backup purposes [10]. The separation of storage from compute resources represents a fundamental architectural pattern in cloud infrastructure, enabling flexible configurations where storage volumes can be detached from terminated instances and reattached to replacement instances during recovery scenarios. EBS volumes are replicated automatically within AWS Availability Zones to provide durability against hardware failures, though this replication occurs within a single datacenter facility rather than across geographically separated locations [10]. The snapshot mechanism in EBS creates point-in-time copies of volumes that are stored in Amazon S3, providing cross-facility durability and enabling volume restoration in different Availability Zones from the original volume location.

EBS snapshots operate through an incremental backup mechanism where only data blocks that have changed since the previous snapshot are transferred to S3 storage [10]. The initial snapshot of a volume captures all allocated blocks, while subsequent snapshots identify and transfer only modified blocks, substantially reducing the storage capacity required for maintaining multiple backup points and minimizing the time required to complete snapshot operations. The incremental nature of EBS snapshots creates dependencies between snapshot generations, where restoration of a specific snapshot requires access to all previous snapshots in the chain back to the initial full snapshot. EBS manages these dependencies automatically, presenting each snapshot as an independent restore point despite the underlying incremental storage architecture [10]. Organizations leveraging EBS snapshots for HANA backup strategies must coordinate snapshot timing with HANA's internal consistency mechanisms to ensure that snapshots capture database states that can be recovered successfully, typically requiring suspension of I/O operations during savepoint execution to guarantee all in-memory modifications have been persisted to storage before the snapshot begins.

Amazon S3 serves as the primary repository for HANA backup files generated through the database's native backup utilities, providing object storage with virtually unlimited capacity that scales automatically as backup data volumes grow over time [12]. S3 implements a distributed storage architecture that automatically replicates objects across multiple storage devices and facilities within an AWS region, delivering high durability characteristics suitable for long-term backup retention. The storage class structure in S3 enables organizations to optimize costs by selecting appropriate storage tiers based on data access patterns and retrieval time requirements [12]. S3 Standard storage provides immediate access to objects with low latency, suitable for recent backups that might be required for restoration operations with minimal recovery time objectives. S3 Standard-Infrequent Access offers reduced storage costs for data accessed less frequently, appropriate for backups transitioning from active status to archival retention where retrieval operations occur infrequently but require rapid access when needed [12].

S3 Glacier storage classes provide further cost optimization for long-term backup retention where retrieval time requirements permit delays measured in minutes to hours rather than immediate access [12]. S3 Glacier Instant Retrieval maintains millisecond access times similar to S3 Standard but at reduced storage costs, targeting backup scenarios where organizations require rapid access to archived data but access frequency remains low. S3 Glacier Flexible Retrieval offers substantial cost savings for backups where retrieval operations can tolerate delays, providing retrieval options ranging from minutes for expedited retrievals to hours for standard retrievals, depending on urgency requirements and cost considerations. S3 Glacier Deep Archive delivers the lowest storage costs for backups requiring extended retention periods where retrieval operations occur rarely and can accommodate delays measured in hours [12]. Organizations implement lifecycle policies in S3 that automatically transition backup objects across storage classes based on object age, enabling automated cost optimization without requiring manual intervention to move aging backups to appropriate storage tiers.

The AWS Backint Agent for SAP HANA provides integration between HANA's backup framework and AWS storage services, enabling direct data flow from HANA processes to S3 without requiring intermediate storage on local file

systems or EBS volumes [6]. The Backint interface represents SAP's standardized API for third-party backup tool integration, defining the protocols through which HANA's backup coordinator communicates with external storage systems during backup and restoration operations. AWS Backint Agent implements this interface to redirect HANA backup streams to S3, managing the complexity of multipart uploads, retry logic for transient failures, and parallel transfer operations that optimize backup throughput [6]. The agent supports configuration options controlling parallelism levels, compression settings, and encryption parameters, enabling organizations to tune backup operations for their specific performance requirements and security policies.

<b>Storage Service</b>	<b>Access Latency</b>	<b>Primary Application</b>	<b>Cost Model</b>
EBS Provisioned IOPS	Sub-millisecond	Active Database Volumes	Premium Tier
EBS Throughput Optimized	Low-moderate	Sequential Backup Operations	Standard Tier
S3 Standard	Milliseconds	Recent Backup Repository	Standard Tier
S3 Standard-IA		Infrequent Backup Access	Reduced Cost
S3 Glacier Instant Retrieval		Archived Rapid Access	Lower Cost
S3 Glacier Flexible Retrieval	Minutes-Hours	Long-term Retention	Economy Tier
S3 Glacier Deep Archive	Hours	Extended Retention	Minimal Cost

Table 1: AWS storage service characteristics for HANA backup operations [10, 12]

**III. DESIGN AND IMPLEMENTATION OF HANA BACKUP STRATEGY ON AWS**

SAP HANA backup operations utilize specialized backup types designed to accommodate the database's in-memory architecture and columnar storage format. The backup mechanisms in HANA must coordinate across multiple database services, including the index server responsible for query processing, the statistics server managing workload metrics, and various specialized servers handling specific functionality depending on HANA configuration [2]. Full data backups capture the complete state of all HANA data volumes at a specific point in time, writing the entire contents of in-memory column tables and row tables to backup destinations. The backup process operates at the persistence layer of HANA's architecture, reading data from savepoint files rather than directly from memory structures, ensuring consistency through HANA's transaction management system that coordinates ongoing modifications during the backup window [2].

The implementation of backup operations in HANA requires careful consideration of the production workload patterns that characterize enterprise database systems. Enterprise database workloads exhibit distinct diurnal patterns with transaction volumes and query complexity varying substantially across different time periods throughout daily operational cycles [7]. Transaction processing workloads typically peak during business hours when user activity reaches maximum levels, while batch processing operations often execute during evening and overnight periods when interactive user activity declines. Organizations must schedule full backup operations during time windows where production workload intensity remains sufficiently low to accommodate the additional I/O load and computational overhead imposed by backup processes without degrading application performance below acceptable

thresholds [7]. The challenge becomes particularly acute for large HANA databases where full backup operations require extended time periods to complete, potentially conflicting with batch processing windows or extending into subsequent peak transaction periods if backup performance does not meet expectations.

Incremental and differential backup strategies in HANA reduce backup windows and storage consumption by capturing only changed data blocks since the previous backup operations. HANA's delta backup mechanism identifies modified data pages within the persistence layer, creating backup sets containing only altered pages rather than complete data volumes [2]. The delta merge operation in HANA continuously consolidates changes from delta storage into main storage structures, creating ongoing data modifications that must be captured by incremental backup processes. Incremental backups form chains where each backup captures changes since the most recent backup of any type, requiring restoration processes to apply the full backup plus all subsequent incremental backups in chronological sequence to reconstruct the database state at a specific point in time [2]. Differential backups capture all changes since the last full backup, regardless of intervening incremental or differential backups, simplifying restoration by requiring only the full backup plus the most recent differential backup without intermediate backup sets.

The coordination of backup operations with HANA's consistency mechanisms ensures that backup images capture database states that can be restored successfully despite concurrent transaction processing during backup execution. HANA employs Multi-Version Concurrency Control to enable simultaneous read and write operations on database tables, maintaining multiple versions of data rows to provide transaction isolation without requiring locks that would block concurrent access [2]. The backup process leverages this versioning mechanism to obtain consistent views of database contents without suspending transaction processing, though backup operations do impose additional I/O load and memory pressure that can affect overall system performance during backup windows. Organizations must monitor backup impact on production workloads through performance metrics including transaction response times, query execution durations, and system resource utilization, to ensure backup operations remain within acceptable performance boundaries.

Data corruption represents a persistent risk in storage systems where hardware failures, software defects, and operational errors can introduce inconsistencies that compromise data integrity. Storage systems experience silent data corruption where bit flips or block errors occur without immediate detection, potentially propagating corrupted data into backup images if validation mechanisms fail to identify corruption before backup execution [8]. The corruption rates vary substantially across different storage technologies and system configurations, with studies demonstrating that corruption detection requires proactive validation mechanisms rather than relying on passive error reporting from storage hardware [8]. HANA implements checksum validation throughout its storage stack to detect corruption during normal operations, and backup processes inherit these validation capabilities to ensure backup images do not contain corrupted data that would prevent successful restoration.

Log backup operations in HANA run continuously in the background, capturing transaction log segments as they fill and writing them to backup storage. The log backup mechanism monitors log segments in the file system, automatically initiating backup operations as segments reach capacity thresholds or time intervals expire [2]. Log backups are essential for point-in-time recovery scenarios where restoration must occur to specific transaction timestamps rather than the time of the most recent data backup. The continuous nature of log backups creates large numbers of small files, with each log segment backup representing an independent object in the backup repository. Organizations configure log backup destinations to write directly to Amazon S3 through AWS Backint Agent, ensuring log backups benefit from S3's durability characteristics and geographic replication capabilities [6]. The log backup frequency balances recovery point objectives against operational overhead, with more frequent log backups reducing potential data loss during disasters but increasing the number of backup files requiring management and the S3 API operations incurred for backup storage.

Backup Strategy	Data Scope	Backup Duration	Storage Consumption	Restoration Dependencie s	AWS Implementa tion
Full Backup	Complete Database	Extended Window	Database Size Equivalent	Single Backup Set	Backint to S3
Incremental Backup	Changed Blocks	Minimal Window	Changed Data Only	Full Plus Chain	Backint to S3
Differential Backup	Cumulative Changes	Moderate Window	Cumulative Since Full	Full Plus Latest	Backint to S3
Log Backup	Transaction Segments	Continuous Background	Log Volume Dependent	Required for Point-in-Time	Backint to S3
EBS Snapshot	Volume Image	Rapid Execution	Incremental Blocks	All Prior Snapshots	AWS Backup Service

Table 2: HANA backup methodology comparison on AWS infrastructure [2, 6, 10]

#### IV. RESTORATION PROCESSES AND RECOVERY TIME OPTIMIZATION ON AWS

Database restoration operations for HANA on AWS must account for substantial data volumes while meeting recovery time objectives that minimize business disruption during disaster scenarios. The restoration process encompasses multiple phases including infrastructure provisioning, data recovery, log replay, and validation testing, before restored databases return to production service [3]. AWS infrastructure enables programmatic provisioning of compute and storage resources through APIs and infrastructure-as-code frameworks, reducing manual effort and potential errors during time-critical recovery operations. Organizations develop automation scripts and templates that codify infrastructure requirements for HANA deployments, enabling consistent recreation of production environments during disaster recovery scenarios [3].

The data recovery phase reads backup files from S3 storage through AWS Backint Agent and reconstructs database structures in the target HANA instance. Large-scale restoration operations benefit from parallel processing where multiple threads simultaneously retrieve different portions of backup data from S3. The distributed architecture of S3 supports high aggregate throughput for parallel read operations, with performance scaling approximately linearly as additional threads access different objects within buckets [12]. Organizations configure Backint parallelism parameters to balance throughput benefits against network bandwidth limitations and memory buffer availability within HANA processes. The restoration throughput depends on multiple factors including the size of individual backup files, the network bandwidth between EC2 instances and S3 service endpoints, and the computational resources available on the EC2 instance for decompression and database reconstruction operations.

Point-in-time recovery operations combine data backup restoration with log backup replay to reconstruct database states at specific transaction timestamps. The recovery process first restores the most recent full backup that precedes the target recovery point, then applies incremental or differential backups as needed based on backup strategy configuration, and finally replays transaction logs up to the specified timestamp [2]. Log replay operations execute sequentially due to transaction ordering dependencies, where each transaction must be applied in the original execution sequence to maintain database consistency. The duration of log replay depends on the transaction volume processed between the data backup timestamp and the target recovery point, making frequent data backups valuable for minimizing log replay duration during recovery operations, even though more frequent data backups consume additional storage capacity and backup window time.

The validation and testing of restored databases encompasses both technical verification of database integrity and functional testing of application operations to ensure complete operational readiness before returning systems to production service. Technical validation includes HANA consistency checks that verify internal data structures, index validation ensuring query operations can access data correctly, and system health checks confirming all database services are operating normally [2]. Application functional testing verifies that business processes execute correctly on the restored database, identifying potential issues including missing or corrupted data that might not be apparent from database-level checks alone. Organizations develop comprehensive testing protocols executed during disaster recovery exercises, refining procedures based on issues encountered during test recoveries to improve success rates during actual disaster scenarios [3].

Cross-region recovery capabilities enable disaster recovery scenarios where primary region failures necessitate database restoration in alternate geographic locations. S3 Cross-Region Replication automatically copies objects from source buckets to destination buckets in different AWS regions, maintaining synchronized backup repositories across geographic locations [12]. The replication process operates asynchronously, with replication lag typically measured in minutes under normal conditions but potentially extending longer during periods of high replication volume or network congestion between regions. Organizations must account for this replication lag when calculating recovery point objectives for cross-region disaster recovery, recognizing that the most recent backups in the primary region may not yet be available in the destination region at the time of disaster declaration. The implementation of cross-region recovery requires preparation including network connectivity configuration, identity and access management policies granting cross-region resource access, and tested recovery procedures accounting for region-specific resource availability and capacity constraints [3].

<b>Recovery Component</b>	<b>Sequential Execution</b>	<b>Parallel Execution</b>	<b>Throughput Enhancement</b>	<b>AWS Optimization Technique</b>
Infrastructure Provisioning	Manual Instance Launch	Automated Template Deployment	Rapid Provisioning	CloudFormation Templates
Volume Restoration	Single Volume Recovery	Multi-volume Snapshots	Concurrent Operations	Fast Snapshot Restore
Data Reconstruction	Single Thread Processing	Multi-thread Backint	Aggregate Bandwidth	S3 Parallel Retrieval
Log Replay	Sequential by Nature	Limited Parallelism	Transaction Dependencies	Frequent Backup Scheduling
Validation Testing	Manual Health Checks	Automated Testing Scripts	Faster Verification	Systems Manager Automation

Table 3: Parallel restoration performance optimization on AWS infrastructure [2, 3, 12]

**V. PERFORMANCE ANALYSIS AND OPTIMIZATION FOR HANA BACKUPS ON AWS**

The performance characteristics of storage systems fundamentally impact backup and restoration operations, with factors including throughput capacity, latency profiles, and concurrent access capabilities determining the time required to complete data protection tasks. Distributed file systems demonstrate architectural patterns for managing large-scale data across multiple storage nodes, employing techniques including data striping to spread individual files across multiple storage devices and parallel I/O operations to aggregate bandwidth from multiple paths simultaneously [9]. These architectural principles apply directly to HANA backup operations on AWS, where backup

data flows from database processes through network infrastructure to distributed S3 storage systems. The parallelization of backup streams across multiple threads enables aggregation of network bandwidth and S3 throughput capacity, substantially improving backup completion times compared to sequential single-threaded approaches.

Data deduplication technologies identify and eliminate redundant data segments within storage systems, providing storage capacity savings particularly valuable for backup repositories containing multiple similar datasets. Deduplication effectiveness varies significantly based on workload characteristics, with structured databases typically exhibiting higher redundancy ratios compared to unstructured data due to repeated occurrences of system metadata, schema definitions, and common data patterns across backup generations [5]. The deduplication process involves chunking data streams into segments, computing cryptographic hashes for each segment to identify content, and maintaining metadata structures that map logical data addresses to physical storage locations containing unique data chunks [5]. Primary storage deduplication systems that operate inline during data writes must balance deduplication computational overhead against storage savings achieved, as hash computations, similarity detection algorithms, and metadata operations consume processing resources that impact overall system performance.

Storage optimization techniques for HANA backups on AWS encompass compression, intelligent storage tiering, and lifecycle management policies that balance performance requirements against cost considerations. Compression algorithms reduce data volumes by exploiting redundancy patterns within data streams, with different algorithms offering varying trade-offs between compression ratios achieved and computational overhead required [5]. HANA's columnar storage format exhibits favorable compression characteristics because data values within columns demonstrate higher similarity compared to heterogeneous row-oriented storage, enabling compression algorithms to achieve substantial size reductions. Organizations configure compression within AWS Backint Agent to reduce data volumes transferred to S3 and stored in backup repositories, though compression does impose computational overhead on EC2 instances that may extend backup windows depending on available CPU capacity and compression algorithm selection [6].

Power management considerations in enterprise storage systems demonstrate optimization opportunities through intelligent workload placement that directs write operations to energy-efficient storage tiers while maintaining acceptable performance levels [11]. These optimization principles translate to cloud environments where organizations are charged based on resource consumption, creating economic incentives for efficiency improvements. Write offloading techniques redirect write-intensive operations including backup workloads, to storage configurations optimized for sequential write patterns, improving both performance and cost efficiency compared to utilizing high-performance storage tiers designed for random access patterns characteristic of transactional database operations [11]. AWS storage service offerings provide hierarchical tiers with varying performance and cost characteristics, enabling organizations to optimize backup storage placement based on access patterns, with recent backups potentially requiring rapid access residing in higher-performance tiers while aging backups transition to lower-cost archival storage as retrieval likelihood diminishes.

The monitoring and optimization of backup operations requires comprehensive visibility into performance metrics, resource utilization patterns, and operational outcomes across backup and restoration processes. Organizations establish baseline performance characteristics during initial HANA deployments, documenting backup completion times, throughput rates, and resource consumption under normal operational conditions [3]. Continuous monitoring identifies performance degradation that might indicate infrastructure issues, configuration problems, or capacity constraints requiring remediation before backup failures occur. Automated alerting mechanisms notify operations teams when performance metrics exceed threshold values, enabling proactive investigation and corrective action. The iterative refinement of backup strategies based on operational experience and monitoring data ensures backup processes remain effective as database sizes grow, workload patterns evolve, and business requirements change over time [3].

Optimization Technique	Storage Efficiency	Backup Window Impact	Processing Overhead	Implementation Phase	Cost Optimization
Compression	Significant Reduction	Moderate Extension	CPU Intensive	During Backup	Reduced Transfer and Storage
Deduplication	High for Structured Data	Minimal Impact	Hash Computation	Post-process or Inline	Storage Capacity Savings
Incremental Backup	Minimal Per Backup	Reduced Duration	Low Overhead	Scheduled Backup	Ongoing Efficiency
S3 Intelligent Tiering	Automatic Optimization	No Impact	None	Lifecycle Policy	Automated Cost Reduction
Lifecycle Policies	Scheduled Transitions	No Impact	None	Policy Configuration	Predictable Savings
Cross-Region Replication	Geographic Redundancy	No Backup Impact	Replication Overhead	Continuous Background	Disaster Recovery

Table 4: Storage optimization techniques and performance trade-offs for HANA on AWS [5, 6, 11, 12]

### CONCLUSION

The deployment of comprehensive backup and restoration strategies for SAP HANA databases on AWS infrastructure requires integration of HANA-native backup capabilities with AWS storage services, automation frameworks, and infrastructure provisioning tools to deliver robust data protection while optimizing cloud resource consumption. Cloud computing has fundamentally transformed the economics and operational models of enterprise IT infrastructure, enabling organizations to provision resources dynamically based on actual demand rather than peak capacity planning, reducing capital expenditure requirements while improving resource utilization efficiency. AWS provides comprehensive ecosystems specifically architected to support enterprise SAP workloads, with compute instances certified for HANA deployments, storage services delivering appropriate performance and durability characteristics, and management tools enabling automation of operational tasks. The architectural flexibility of AWS enables organizations to design backup strategies tailored to specific operational requirements, balancing factors including backup window constraints, recovery time objectives, storage cost optimization, and regulatory compliance mandates. Organizations implementing HANA on AWS benefit from the separation of compute and storage resources, enabling independent scaling of backup infrastructure, the geographic distribution of AWS regions supporting cross-region disaster recovery architectures, and the integration of AWS services providing centralized management of backup operations across multiple HANA systems. The evolution of backup technologies continues with emerging capabilities, including machine learning algorithms for predictive optimization, automated anomaly detection, identifying potential failures before they impact operations, and intelligent recovery orchestration, selecting optimal restoration approaches based on specific scenario requirements. Organizations planning HANA migrations to AWS should conduct thorough assessments of backup requirements during planning phases,

implement comprehensive testing of backup and restoration procedures, including full disaster recovery exercises validating cross-region recovery capabilities, and establish ongoing optimization processes continuously refining backup strategies to align with evolving business requirements and AWS capability enhancements.

### REFERENCES

- [1] Michael Armbrust et al., "A view of cloud computing," *Communications of the ACM*, 2010. doi: 10.1145/1721654.1721672. Available: <https://cacm.acm.org/magazines/2010/4/81493-a-view-of-cloudcomputing/>
- [2] SAP, "SAP HANA Administration Guide for SAP HANA Platform," 2022. Available: [https://help.sap.com/doc/eb75509abofd1014a2c6ba9b6d252832/2.0.05/enUS/SAP\\_HANA\\_Administrati on\\_Guide\\_en.pdf](https://help.sap.com/doc/eb75509abofd1014a2c6ba9b6d252832/2.0.05/enUS/SAP_HANA_Administrati on_Guide_en.pdf)
- [3] AWS Amazon, "SAP HANA on AWS Operations Guide,". Available: <https://docs.aws.amazon.com/sap/latest/sap-hana/sap-hana-on-aws-operations.html>
- [4] Ang Li et al., "CloudCmp: Comparing public cloud providers," *IMC '10: Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*, 2010. doi: 10.1145/1879141.1879143. Available: <https://dl.acm.org/doi/10.1145/1879141.1879143>
- [5] Dutch T. Meyer and William J. Bolosky, "A study of practical deduplication," *ACM Transactions on Storage (TOS)*, Volume 7, Issue 4, 2012. doi: 10.1145/2078861.2078864. Available: <https://dl.acm.org/doi/10.1145/2078861.2078864>
- [6] AWS Amazon, "AWS Backint Agent for SAP HANA,". Available: <https://docs.aws.amazon.com/sap/latest/saphana/aws-backint-agent-sap-hana.html>
- [7] W. W. Hsu et al., "Characteristics of production database workloads and the TPC benchmarks," *IBM Systems Journal*, Volume 40, Issue 3, 2001. Available: <https://ieeexplore.ieee.org/document/5386932>
- [8] Lakshmi N. Bairavasundaram et al., "An analysis of data corruption in the storage stack," *ACM Transactions on Storage (TOS)*, Volume 4, Issue 3, 2008. doi: 10.1145/1416944.1416947. Available: <https://dl.acm.org/doi/10.1145/1416944.1416947>
- [9] Sanjay Ghemawat et al., "The Google file system," *SOSP '03: Proceedings of the nineteenth ACM symposium on Operating systems principles*, 2003. doi: 10.1145/945445.945450. Available: <https://dl.acm.org/doi/10.1145/945445.945450>
- [10] Amazon, "Amazon EBS snapshots,". Available: <https://docs.aws.amazon.com/ebs/latest/userguide/ebssnapshots.html>
- [11] Dushyanth Narayanan et al., "Write Off-Loading: Practical Power Management for Enterprise Storage," *FAST '08: 6th USENIX Conference on File and Storage Technologies*. Available: [https://www.usenix.org/legacy/event/fast08/tech/full\\_papers/narayanan/narayanan.pdf](https://www.usenix.org/legacy/event/fast08/tech/full_papers/narayanan/narayanan.pdf)
- [12] Amazon, "Amazon S3 Storage Classes,". Available: <https://aws.amazon.com/s3/storage-classes/>