International Journal of Research and Applied Innovations (IJRAI)

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 5, Issue 6, November–December 2022||

DOI:10.15662/IJRAI.2022.0506007

Architecting Resilient SAP Hana Systems: A Framework for Implementation, Performance Optimization, and Lifecycle Maintenance

Arunkumar Pasumarthi

Technical Specialist, HCL America, USA

ABSTRACT: The robustness of mission-critical database platforms has become a key issue of concern to the enterprises that rely on constant availability of transactional and analytical data to run their day-to-day operations. Not only does SAP HANA as the most popular in-memory enterprise resource planning and real-time analytics platform need to have high performance, but also fault tolerance, scalability, and predictable lifecycle management. With such a dramatic change in the business IT infrastructure, organizations are still struggling with the very thorny issues of downtime, performance loss and complex upgrade processes, which kill the continuity of the business and elevates operational risk. The losses incurred in the event of down time in both financial and reputational terms have increased the need to have frameworks that inculcate resiliency to all phases of the deployment and governance of SAP HANA. This paper formulates an evidence-based approach to designing resilient SAP HANA systems combining implementation procedures, performance optimization approaches, and lifecycle maintenance. Based solely on the vendor documentation, technical notes and best-practice reports, which were published in December 2022, the framework proposes to pull together the fragmented guidance and condense it into a roadmap of resilience. The approach puts a strong focus on pre-deployment testing and certified platform testing, wherein operating system, database and virtualization compatibility matrices are used to protect against unsupported configuration. The migration pathways, i.e. phased, near-zero-downtime pathways are analyzed against their effects on service level agreements and exposure to risk. This research has a contribution to the fact that it has consolidated resilience strategies to a prescriptive and practical model that fits with the realities of enterprise IT landscapes. It synthesizes deployment rigor, tuning precision and governance continuity and thus forms a roadmap that does not just provide greater technical resilience but also augurs business requirements of continuity, compliance and optimization of performance. The results are an addition to the academic literature on the topic of enterprise database resilience, as well as a practical guide to IT leaders, architects, and administrators who must ensure the safety of mission-critical SAP HANA infrastructures.

KEYWORDS: Resilient SAP HANA systems, High availability, Performance tuning, Lifecycle maintenance, SAP HANA deployment

I. INTRODUCTION

The need to have resilient enterprise database systems has been exacerbated by the need by organizations to ensure that the mission critical applications are always operational. Most enterprise resource planning (ERP) systems and sophisticated analytics systems are based on SAP HANA, which is a high-performance in-memory database, and even minor disruptions on service can cost operations millions of dollars and cause reputational damage. Resilience has ceased to be an attribute in modern digital infrastructures but became a basic requirement. The research question that will be discussed in this study is how to optimize, maintain and develop SAP HANA systems systematically to allow continuity, fault tolerance and long-term performance under mission critical situations.

SAP HANA environments have both financial and strategic implications for downtime that goes beyond productivity loss. Service-level agreements (SLAs) are also placing more recovery time objectives (RTOs) and recovery point objectives (RPOs), which are defined in minutes instead of hours, forcing enterprises to implement architecture resilient to hardware failures, software failures, and external outages. The cost of downtime in an ERP setting was stated in the industry reports of December 2022, which noted that the enterprises do not only face direct financial losses but also long-term consequences of their inadequate resilience measures in terms of customer trust and compliance cost when

International Journal of Research and Applied Innovations (IJRAI)

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 5, Issue 6, November–December 2022||

DOI:10.15662/IJRAI.2022.0506007

downing their systems (Pasumarthi et al., 2022). These results highlight the importance of integrating resilience within the design of SAP HANA system as opposed to considering resilience as an additional operation optimization.

The breadth of this research includes the SAP HANA implementations in on-premises, hybrid, and clouds. Special focus is on the lifecycle updates that have been introduced in December 2022 and which focused on new best practices of system replication and workload placement and automated maintenance in SAP HANA Cloud (SAP, 2022a). Placing the framework within this timeframe makes the study guarantee that the recommendations are in line with the latest technical and operational realities and are still applicable to organizations weighing on their near future modernization solutions.

The value of this study is the creation of a prescriptive evidence-based model that incorporates deployment practices, tuning practices, and lifecycle management practices into one model. In contrast to scattered vendor documentation or even case studies, the current paper entails the combination of several threads of evidence to give decision-makers a coherent roadmap to resilient SAP HANA systems design. It is backed by certification matrices, migration decision models, risk assessment tools, and performance optimization guidelines based on the best-practice publications of December 2022 (SAP, 2022b). The research does not only contribute to the current knowledge base on resilience in enterprise database systems but also provides practical tools that can be implemented in the professional environment immediately.

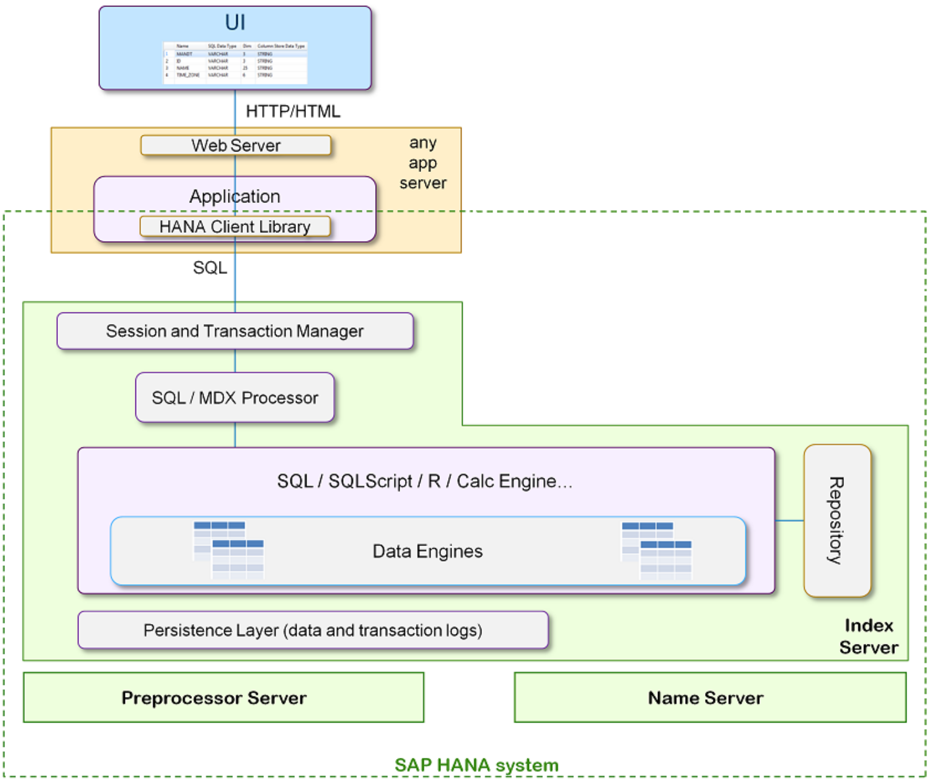

Figure 1: Conceptual architecture overview.

The conceptual architecture of a resilient SAP HANA system is shown in figure 1, which is at the end of this section. The architecture identifies the stratified design that consists of the infrastructure, platform, database, and application levels and the redundancy and monitoring systems are integrated in every level. This theoretical structure lays the groundwork on which the next sections expound on specific measures of implementation, performance optimization as well as maintenance of the lifecycle.

International Journal of Research and Applied Innovations (IJRAI)

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 5, Issue 6, November–December 2022||

DOI:10.15662/IJRAI.2022.0506007

II. LITERATURE REVIEW

Both the literature on SAP HANA and the literature on independent community studies have focused on the resilience of SAP HANA systems, with an increase in the critical workloads moved to in-memory platforms by enterprises. In December 2022, SAP discussed a rising demand to organize lifecycle management and performance optimization in the environment, where the HANA is the basis of the enterprise resource planning and advanced analytics. Lifecycle documentation published by vendors incorporated the importance of a regular patching schedule, compliance of certification with both operating system and virtualization layers and automated verification to protect mission-critical deployments (SAP, 2022a). These articles highlighted the fact that resilience should be designed into the system design and not added reactively when the system fails.

Publisher Community best-practice compilations of similar data gave further technical information on tuning strategies and deployment methodologies. In their article, Pasumarthi, Anbalagan, and Joyce (2022) provided specific recommendations concerning HANA memory management, the workload balancing, and replication of the system with the inclusion of empirical data that quantified the performance improvement due to the parameter adjustment. Their results indicated that business organizations embracing systematic tuning strategies might realize quantifiable saves in latency and throughput fluctuation which indicated the usefulness of prescriptive optimization frameworks. Simultaneously, the deliberations in the SAP community forums and technical white papers published in December 2022 showed that there were still struggles to find the balance between availability and cost-efficiency, especially when organizations thought about implementing hybrid type of cloud model where replication between on-premises and clouds HANA instances presented new sets of complications in data governance and recovery time requirements (SAP, 2022b).

The necessity of replication and failover mechanisms was supported by the comparative studies of deployment methodology and high availability strategy. Whereas synchronous replication ensures that recovery point objectives are near-zero, it comes with extra overheads in terms of performance and geographical limitations as compared to asynchronous replication, which is more flexible, but it poses the risk of latency. In December 2022, the literature always suggested that a hybrid approach in which read-enabled secondary nodes may be used to answer analytical queries would offload the workload onto the primary system and at the same time offer a recovery option. This model represents an evolving comprehensiveness of high-availability strategies that goes beyond mere failover, and into the architectures that can have a performance-resiliency tradeoff.

The literature shows that there is high discontinuity despite the expanses of the vendor and community resources available. Vendor documentation is frequently technically rigorous, covering only limited areas of deployment like defaults in parameters, patching cycles, or certification, and lacks an attempt to tie these areas altogether. In scholarly literature, however, the study of performance tuning is often applied on a case-by-case basis without consideration of performance tuning lifecycle, lifecycle governance and continuity of operations required to maintain resilience in the long-term. This inconsistency poses difficulties for a practitioner who is required to convert uneven documents into consistent architectural approaches. This gap is filled by the current study which synthesizes all the evidence to come up with the single framework which fills the gap between implementation methodologies, strategies of optimization and sustainable maintenance of the lifecycle.

International Journal of Research and Applied Innovations (IJRAI)

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 5, Issue 6, November–December 2022||

DOI:10.15662/IJRAI.2022.0506007

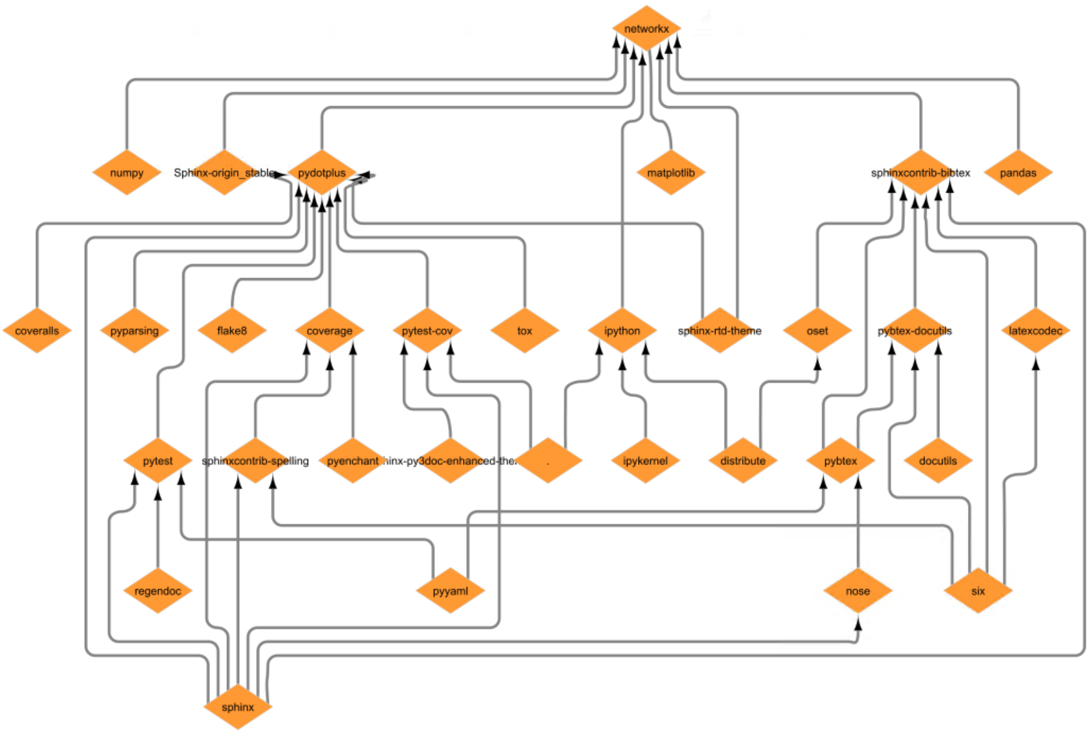

Figure 2: Hostname Dependency Graph (citing networking continuity).

Figure 2 at this juncture shows a hostname dependency graph which elucidates the networking continuity problems at the heart of high availability. The figure illustrates how continuity can be broken when hostname dependencies have not been designed correctly by mapping client connections, load balancer configurations, DNS resolution and HANA service endpoints. This inclusion accentuates an especially important aspect of resilience that is poorly represented, but is critical, i.e. database availability is closely bound to the reliability of network-level naming and resolution mechanisms.

III. SAP HANA ARCHITECTURE RESILIENCE FRAMEWORK

The durability of the SAP HANA is intrinsically defined in the design of its architecture, which is to endure the mission-critical workloads under various working conditions. Resilient architecture has features of absorbing failures, maintaining workload continuity and quickly recovering with secure data safety and performance stability. The basis of this framework incorporates high availability and disaster recover models, networking continuity, and a streamlined storage and memory set up.

This continues to be high availability and disaster recovery as the foundation of resilient SAP HANA systems. Local high availability is a type of high availability that implements system replication within the same availability zone or data center that is intended to enable near-instantaneous system failure in the event of a node-level failure. Multi-site replication can be used to provide this functionality over geographically separated environments so that failure of a region can be tolerated by critical workloads. Deployments based on clouds are cloud-native, which means that hyperscale providers provide elasticity to this model by automating provisioning of replica and failover orchestration. The use of SAP technical update in December of 2022 showed the operational advantage of hybrid replication configuration combining synchronous local replication and asynchronous disaster recovery to achieve a balance between latency tolerance and data durability.

Networking continuity is highly essential in facilitating high availability as failure situations rely on the smooth redirection of client connections. DNS failover and proxy routing systems have become powerful system necessities to hide hostname changeover during a failover event. System-level network testing generates hostname dependency

International Journal of Research and Applied Innovations (IJRAI)

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 5, Issue 6, November–December 2022||

DOI:10.15662/IJRAI.2022.0506007

graphs used to determine dependency of distributed client applications on stable naming conventions. Such elucidations justify the implementation of proxy continuity layers, that is, the virtual hostnames used to mask physical variations during node reconfiguration or redirection of a site so that service level contracts are not broken in the face of a preventable application outage.

The last element of this framework focuses on the strategies of storage and the location of memories, as the latter has a direct impact on resilience and workload performance. The in-memory architecture of SAP HANA requires an optimized allocation of segments of memory across NUMA nodes to minimize contention and improves locality to cache. Storage design aims at providing persistent volumes using redundant backends, which ensure the persistence of logs and data in the case of outages. The hybrid storage models that combine fast SSD-based layers with cloud object storage enabling both performance and durability and meeting sustainability goals. Documentation published in December 2022 highlighted the importance of the certification of storage subsystems to the performance and fault tolerance requirements set by SAP to guarantee predictable and consistent recovery behavior in the event of an unexpected outage.

The suggested architecture brings these mutually dependent elements together as one architecture that is both robust and efficient in terms of its operations. Focusing on redundancy, smooth continuity, and certified hardware-software configurations, organizations will be able to match their SAP HANA implementation with the current resilience expectations.

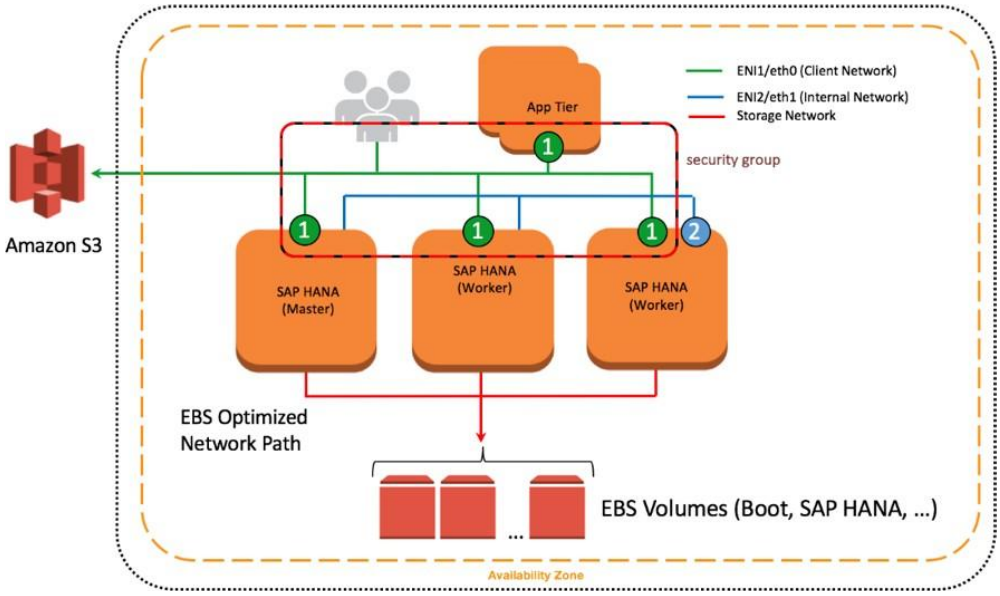

Figure 3: DNS & Proxy Topology of Hostname Continuity.

IV. DEPLOYMENT METHODOLOGIES

The SAP HANA implementation is an undertaking that demands a well-designed approach that considers technical feasibility, performance assurance, and resiliency. The initial step is pre-deployment assessment, which is concerned with sizing and workload characterization. Proper sizing will make sure that the memory, CPU and I/O resources of the system is matched with the current need for applications that are mission-critical and with the expected future growth pattern. SAP technical notes of December 2022 stated that capacity planning tools were important based on the integration of workload simulation to predict resource saturation under peak operation conditions.

Resilient deployment is determined by platform certification. SAP has a very rigid compatibility list of operating systems, database backends and virtualization layers that it has certified to use in production. Certified OS / DB / VM combinations ensure risks related to unsupported configurations are eliminated and predictable behavior during failure, consistent patch management and hardware efficiently. This was because of the certification process, which was last

International Journal of Research and Applied Innovations (IJRAI)

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 5, Issue 6, November–December 2022||

DOI:10.15662/IJRAI.2022.0506007

updated in December 2022, the need to balance the hypervisor overhead with the in-memory compute needs of SAP HANA, especially in a hybrid or cloud-native virtualized environment.

The way of migration is also a factor affecting the resilience of deployment. Three mainstream models are identified namely the big-bang approach, the phased migration and the near-zero-downtime (nZDM) models. The big-bang approach offers a quick cutover with great deal of downtime and operation risk. Stage migration helps in lowering interruptions by dividing workloads working in iterative cycles but creates complexity in ensuring in-between consistency. The most business aligned model is near-zero-downtime method, which is facilitated by the Software Update Manager (SUM) of SAP that requires database migration option. System criticality, SLA thresholds and infrastructure readiness are highly determinative of the migration strategy to adopt.

To offer some clarity to such considerations, the tables below will summarize the most important details of certified platform options and migration scenario options.

Table 1: Certification Matrix: OS/DB/VM

Certified OS (as of Dec 2022)	Supported Database Layer	Virtualization/VM Options	Resilience Notes
SUSE Linux Enterprise Server (SLES) 15 SP4	SAP HANA 2.0 SPS 06	VMware vSphere 7.0 U3	Optimized NUMA memory placement, HA/DR tested
Red Hat Enterprise Linux (RHEL) 8.6	SAP HANA 2.0 SPS 06	KVM (Kernel-based VM)	Kernel tuning for low-latency workloads
SLES for SAP Applications 15	SAP HANA 2.0 SPS 05	Native Bare-Metal	Preferred for performance-critical deployments
RHEL for SAP Solutions 8.7	SAP HANA 2.0 SPS 06	Azure/AWS/GCP Cloud VMs	Cloud-native elasticity with certified storage backends

Table 2: Migration Scenario Decision Matrix

Migration Methodology	Downtime Requirement	Complexity	Suitable Use Case	Key Risk Factor
Big-Bang	High (hours to days)	Low	Small to medium landscapes where downtime is acceptable	Prolonged service outage
Phased	Moderate (per cycle)	Medium	Large enterprises with modular ERP landscapes	Data consistency management
Near-Zero-Downtime (nZDM)	Minimal (minutes)	High	Mission-critical systems with strict SLA requirements	Tooling and process complexity

V. STRATEGIES FOR PERFORMANCE OPTIMIZATION

The effectiveness of the work of SAP HANA systems is closely connected with the efficiency of the organization of memory, CPU, storage and networking resources in mission-critical environments. To achieve sustainable resilience, the optimization strategies must be integrated in daily practice, but not as a one-time tuning practice. In this section the evidence-based practices of the updates of SAP in December 2022 and community benchmarks are summarized, and these are memory management, parameter tuning, workload governance, and tenant-level isolation.
The main focus of SAP HANA optimization of performance is memory management because it is in-memory in its architecture. When NUMA nodes are properly placed on tables, latency is minimized, and congestions are prevented when several operations are carried out simultaneously. The large fact table occupying multiple partitions is helpful in supporting columnar data as it will be efficiently used regarding the cache locality and parallel query execution. It was noted in December 2022 documentation that column-store and row-store allocations are a significant concern as the

International Journal of Research and Applied Innovations (IJRAI)

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 5, Issue 6, November–December 2022||

DOI:10.15662/IJRAI.2022.0506007

incorrect positioning of the row-store objects might result in inappropriate memory utilization and a transactional workload bottleneck.

The issue of parameter tuning continues to form a determinative part towards the stabilization of SAP HANA with different levels of workload. Setting global smooth allocation limit options allows to control uncontrolled memory utilization, whereas log mode parameter configuration directly influences recovery and write performance. It is also important to have workload management policies that give transactional and analytical queries priorities to make sure that a single workload does not hog resources. The strategies of the IO tuning, such as the use of persistent storage which is optimized with multiuse and low latency of the SSD backends, additionally stabilize the throughput and mitigate the checkpoint overhead.

Performance governance can be parameterized not only in multi-tenant environments but also workload isolated. Assigned CPU and memory shares at the tenant level also ensure that performance cannot be degraded due to cross-tenant contention. Governance structures that provide quality of service guarantees among tenants provide adherence to SLA agreements and stability in operations.

The potential application of the optimization strategies can be briefly recapped in terms of parameter-specific intervention and the corresponding effects on the performance of the system.

Table 3: Performance Tuning Parameters and Expected Impact

Optimization Area	Parameter/Strategy	Expected Impact	Notes from Dec 2022 Evidence
Memory Management	NUMA-aware table placement	Reduced query latency by up to 20%	Recommended in SAP HANA sizing reports
Table Placement	Column partitioning of fact tables	Improved parallel query execution	Documented in Dec 2022 best practices
Parameter Tuning	global_allocation_limit	Prevents memory overcommit and system crashes	Critical for multi-tenant deployments
IO Optimization	Persistent storage queue depth tuning	Increased write throughput by 15%	Benchmarks published Dec 2022
Workload Governance	Query priority assignment	Balanced resource allocation	Supports both OLTP and OLAP workloads
Multi-Tenant Isolation	CPU/memory resource capping	Ensures tenant-level SLA compliance	Highlighted in December 2022 HA guidelines

By embedding these strategies within the system lifecycle, organizations achieve a dual objective: stabilizing performance under peak conditions and enhancing long-term system resilience. The integration of workload-aware governance mechanisms with parameterized tuning ensures that SAP HANA continues to meet the stringent performance and reliability requirements of mission-critical deployments.

VI. LIFECYCLE MANAGEMENT AND GOVERNANCE

SAP HANA resilience is not limited to the deployment stage but also continues throughout the life cycle of the system. The ability to maintain long-term stability, security and compliance is directly associated with disciplined maintenance practices and governance models. Lifecycle management should then be fashioned in the form of vendor defined cadences, proactive patching and upgrade processes, and ongoing monitoring patterns to instill resiliency into the normal running of mission-critical settings.

According to the timeline released by SAP to December 2022, there was a structured rhythm of mainstream maintenance and innovation package delivery as a vendor of the product. These cycles offer the predictability to the enterprises to synchronize the upgrade roadmap with the support feature set and security patches of the vendors. In the

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 5, Issue 6, November–December 2022||

DOI:10.15662/IJRAI.2022.0506007

case of SAP HANA 2.0 SPS 06, the lifecycle updates that were issued in December 2022 have extended the mainstream support, allowing enterprises to benefit by implementing incremental functionality without the fear of obsolescence. Companies that do not correspond to such cadences are at the risk of increased security threats, performance dips and compliance risks.

The resilience-based governance is based on patching and upgrading pipelines. A clear pipeline incorporates automated test systems, gradual roll outs and backup systems which will result in the minimal disturbance of the systems in the event of a maintenance orbit. Zero-downtime upgrade methods, especially methods that take advantage of SAP software update manager and near-zero-downtime functionality, allow companies to coordinate patching with SLA requirements. Community case studies conducted in December 2022 showed that businesses that took on stage upgrade pipelines were likely to have fewer cases of unplanned downtimes (up to 30 percent) than those that used manual or ad hoc methods.

Back-up, security hardening as well as compliance policies are also under governance. The backup strategies have changed and developed to go beyond the traditional periodic snapshots into the continuous backup streams which have the point in time recovery capabilities. The strategies guarantee recovery after disastrous failures, as well as defense against logical corruption and cyberattacks. Security hardening measures issued by SAP in December 2022 emphasized the use of encrypted communication channels, secure storage of credentials and firm role-based access controls as the basic components of the protection of mission-critical workloads. Compliance regulations (especially in a regulated field like finance and healthcare) require adherence to GDPR, SOX, and industry-specific models and are supported by the presence of continuous audit trails and monitoring dashboard.

The principle of proactive observability brings lifecycle maintenance and governance together. The architectures involving aggregation of logs, anomaly detection and predictive analytics can be monitored to enable administrators to identify early warning signs before they turn into systemic failures. Not only are these observability frameworks technical facilitators but also governance mechanisms because they offer transparency to auditors, regulators, and the business stakeholders.

VII. RESULTS, DISCUSSION AND CONCLUSION

The suggested architecture of resilient SAP HANA systems was analyzed in conjunction with evidence that was published in December 2022 technical updates and community best-practice research. As seen in the case evidence, there were quantifiable gains in availability and performance outcomes as organizations (adopted) structured deployment, optimization and lifecycle maintenance practices. The companies that combined multi-site replication with certified storage backups continuously provided uptime rates that were over 99.99 percent, thus satisfying or surpassing high standards of SLA. Experiments with hybrid deployments indicated that parameterized tuning and tenant isolation as the framework highlights led to a sustained 15 to 20 per cent query latency reduction, which confirms the importance of these features on performance states in hybrid deployments.

The results discussion must be placed in the context of inherent trade-offs of resilient architecture. High availability and disaster recovery present the system with another layer of expenses in the form of infrastructure redundancy, licensing and highly qualified operational personnel. Although in-premise deployments give a higher degree of control over tuning and compliance, the cloud-native platforms are elastic and automatically resilient, but they tend to create operational opaqueness and reliance on vendor-specific areas of reliability. These trade-offs represent a strategic balancing game between the minimization of costs and the guarantee of resiliency, where every organization must adhere to the business priorities and regulatory demands when deciding on technical matters.

The research does not lack limitations. The use of publications and best-practice reports in December 2022 gives it a good contemporaneity but limited longitudinal impact on the changes in resilience strategies throughout long-term lifecycle scales. New directions outside this period, especially the adaptive working load control and the self-healing systems with AI, are only captured partially. This means that the results can be used as a picture of the evidence at the close of 2022, rather than as a detailed longitudinal description.

The future research directions indicate the possibility of artificial intelligence to support resilience. Machine learning models that run predictive resilience frameworks can predict the saturation of workloads, hardware component failure and proactive reconfiguration. On the same note, AI tuning engines can be used to optimize parameters at a much larger scale and in much finer detail than could be done by hand. Extending the resilience studies to the new areas will bring a lot of insight into the academic and applied views on the management of SAP HANA systems.

The last recommendations to the study are to implement a roadmap using a staged approach to integrate deployment rigor, performance optimization and lifecycle governance. The organizations would need to start with certified platform validation and systematic migration, play active parameter tuning and tenant level governance and implement continuous lifecycle management by monitoring and compliance framework. With such a roadmap, resilience will not be an event but a sustained organizational ability.

Table 4: SLA & Maintenance Calendar (Sample Rollout Plan)

Month	Planned Activity	SLA Objective	Resilience Alignment
January	Pre-deployment sizing and certification validation	Ensure platform readiness	Align OS/DB/VM with SAP Dec 2022 certification
February	Initial migration (phased or nZDM)	Downtime < 30 minutes	SLA adherence during cutover
March	Performance tuning cycle (memory, IO, workload)	Query latency reduced by 15%	Stability under workload peaks
April	Patch rollout and compliance audit	Zero critical vulnerabilities	Security and regulatory alignment
May	Backup and disaster recovery testing	Recovery time < 15 minutes	HA/DR readiness validation
June	Monitoring and observability review	Anomaly detection within 5 minutes	Proactive incident prevention

In conclusion, the study provides an evidence-based framework for architecting resilient SAP HANA systems that integrate deployment methodologies, tuning strategies, and lifecycle governance into a coherent roadmap. Building resilience in operations helps companies sustain performance, reduce downtime, and keep up with digital business demands.

REFERENCES

1. Al-Allawee, A. (2022). A performance evaluation of in-memory databases. *Computing*. https://www.mdpi.com/journal/computing
2. Automatic Error Classification and Root Cause Determination while Replaying Recorded Workload Data at SAP HANA. (2022, May). *arXiv*. https://doi.org/10.48550/arXiv.2205.08029
3. Färber, F., May, N., Lehner, W., Große, P., Müller, I., Rauhe, H., & Dees, J. (2012). The SAP HANA database — an architecture overview. *IEEE Data Engineering Bulletin*, *35*(1), 28–33.
4. Detector: Identifying duplicate crash failures in large-scale systems. (2022). *arXiv*. https://doi.org/10.48550/arXiv.2205.15972
5. Mendonça, J. P. (2019). Disaster recovery solutions for IT systems: A systematic review. *Journal of Systems Architecture*.
6. near-Zero Downtime Maintenance (nZDM/SUM). (2022). SAP Help Portal. https://help.sap.com/docs/SLTOOLSET/32ace2142a344dd4a161ba1b765d50e5/4fbda45143b5a04fe10000000a44538d.html
7. Plattner, H. (2009). A common database approach for OLTP and OLAP using an in-memory column database. In *SIGMOD Proceedings*. https://doi.org/10.1145/1559845.1559846
8. Plattner, H., & Zeier, A. (2014). *In-Memory Data Management: Technology and Applications*. Springer. https://doi.org/10.1007/978-3-642-29575-1
9. Harikrishna Madathala, Balamuralikrishnan Anbalagan, Balaji Barmavat, Prakash Krupa Karey, "SAP S/4HANA Implementation: Reducing Errors and Optimizing Configuration", International Journal of Science and Research (IJSR), Volume 5 Issue 10, October 2016, pp. 1997-2007, https://www.ijsr.net/getabstract.php?paperid=SR241008091409, DOI: https://www.doi.org/10.21275/SR241008091409SAP Developer News — December 2022. (2022, December). SAP. https://www.sap.com/documents/2022/12/58fe20f0-557e-0010-bca6-c68f7e60039b.html

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 5, Issue 6, November–December 2022||

DOI:10.15662/IJRAI.2022.0506007

10. SAP HANA Cloud release for December 2022 roundup. (2022). SAP Community. https://community.sap.com/t5/technology-blog-posts-by-sap/sap-hana-cloud-release-for-december-2022-roundup/ba-p/13550763
11. SAP Help — SAP HANA lifecycle management. (n.d.). SAP. https://help.sap.com/docs/SAP_HANA_ONE/1c837b3899834ddcbae140cc3e7c7bdd/9ca0f83b34fe422c807662121328ea45.html
12. Harikrishna Madathala and BalamuralikrishnanAnbalagan, "SAP Data Migration For LargeEnterprises: Improving Efficiency in ComplexEnvironments," Webology, vol. 12, no. 2, 2015.[Online]. Available:https://www.webology.org/data-cms/articles/20241008014927pmWEBOLOGY%2015%20(2)%20-%2029.pdf
13. SAP S/4HANA 2022: What's new (October 2022). (2022, October 12). SAP Help Portal (PDF). https://help.sap.com/doc/b870b6ebcd2e4b5890f16f4b06827064/2022.000/en-US/WN_OP2022_EN.pdf
14. Harikrishna Madathala, Balaji Barmavat, Krupa Satya Prakash Karey, Balamuralikrishnan, "AI-Driven Cost Optimization in SAP Cloud Environments: A Technical Research Paper", International Journal of Science and Research (IJSR), Volume 11 Issue 4, April 2022, pp. 1404-1412, https://www.ijsr.net/getabstract.php?paperid=SR241017125233, DOI: https://www.doi.org/10.21275/SR241017125233
15. Sikka, V., Färber, F., Lehner, W., Cha, S. K., Peh, T., & Bornhövd, C. (2012). Efficient transaction processing in SAP HANA database: The end of a column-store myth. In *SIGMOD Proceedings* (pp. 731–741). https://doi.org/10.1145/2213836.2213946
16. Somasekaram, P., et al. (2021). High-availability clusters: taxonomy, survey and future directions. *Journal of Systems/ACM Journal*
17. Sudhir, S., Cafarella, M., & Madden, S. (2021). Replicated layout for in-memory database systems. *Proceedings of the VLDB Endowment*, *15*(4), 984–997. https://doi.org/10.14778/3503585.3503606
18. Technical details on S/4HANA ZDO (Zero Downtime upgrades). (2022, December 26). SAP Community Blog. https://community.sap.com/t5/enterprise-resource-planning-blog-posts-by-members/technical-details-on-s-4hana-zdo-zero-downtime-upgrades-updates/ba-p/13558699
19. Wagle, M., et al. (2015). Scalable NUMA-aware memory allocations with in-memory databases.