



Engineering Fail-Safe SAP Hana Operations in Enterprise Landscapes: How SUSE Extends Its Advanced High-Availability Framework to Deliver Seamless System Resilience, Automated Failover, and Continuous Business Continuity

Sankar Thambireddy

Senior Technology Consultant, SAP America Inc., USA

Venkata Ramana Reddy Bussu

Senior Cloud Solutions Engineer, CodeTech Inc (DTE Energy), USA

Arunkumar Pasumarthi

Technical Specialist, HCL America, USA

ABSTRACT: This paper offers a detailed analysis of how SAP becomes strong in ensuring the reliability of the operations of SAP HANA in enterprise systems with the support of SUSE advanced high-availability (HA) framework. Since SAP HANA remains a mission-critical in-memory database solution, business continuity, speedy recovery, and database consistency have become crucial to organizations that operate within highly competitive digital ecosystems. SUSE responds to these needs and adds value to the features of the native SAP HANA with robust clustering, automated failover, and management of resources, fully minimizing downtime and protecting business continuity. The paper explores the HA architecture at SUSE, the agents of the resources, and fencing mechanisms and how they are combined with SAP HANA system replication in both scale-up and scale-out deployments. Through a comparative assessment paradigm by assessing recovery time objective (RTO), recovery point objective (RPO), failover time, and system availability, the study compares native SAP HANA resilience tools with the enhanced structure of SUSE. The results of the 2022 SUSE documentation and cloud-based implementation on AWS, Google Cloud, and Microsoft Azure show that with SUSE, failover speeds are much lower and operations more resilient but come with additional configuration complexity and overhead. The results show practical trade-offs, performance, and enterprise-wide advantages to implementing SUSE HA and provide practical advice to IT managers and system architects who need to make sure they maintain ongoing operations. In addition, the paper presents the existing limitations in containerized and cloud-native HANA systems and gives guidelines and further research directions on the development of high-availability strategies in present-day enterprise environments.

KEYWORDS: SAP HANA System Replication, SUSE High Availability Extension, Automated Failover, System Resilience, Business Continuity

I. INTRODUCTION

SAP HANA has emerged as one of the most essential in-memory databases in the contemporary enterprise environment, which drives real-time analytics and transactional processing to global organizations. With the expansion of businesses into multi- and hybrid-cloud environments, the need to develop failure resilient mechanisms to protect the availability of data and guarantee business continuity has grown. System resilience in SAP HANA environments has become so important because, even temporary outages can lead to a significant loss of revenue, inconsistency of data, and a lack of continuity in business processes, which is especially important in such industries as finance, telecommunications, and healthcare (SUSE, 2022a).

Although native high-availability capabilities exist in SAP HANA, such as system replication and automatic host failure, those are often faced with issues that reveal the shortcomings of those features by enterprises. Difficulties in design of system replication, the possibility of split-brain, and lack of complete automated mechanism of failover in



distributed landscapes provide vulnerabilities that jeopardize service-level agreements. Replication, detection, and takeover should also be orchestrated in an automated manner as companies implement increasingly geographically distributed architecture, and native SAP HANA tools are incapable of doing this (Google Cloud, 2022).

This study will deal with the issue of developing and engineering a fail-safe SAP HANA process capable of surviving failures of the system and site without human interference. The role of SUSE in expanding the high-availability framework through its sophisticated resource agents, cluster automation, and HA/DR provider connections are the subjects of this study, in which businesses can implement seamless resilience and total continuity in their business. Complemented with High Availability Extension, SAP Applications by SUSE Linux Enterprise Server allow an automated failover, resource management, and reduced recovery times to bridge a crucial gap that has been left by standard SAP mechanisms (SUSE, 2022b).

This paper aims at critically assessing the contribution of the improvements at SUSE towards automated resilience in SAP HANA deployments. The study quantifies the benefits of SUSE-enhanced clusters over un-enhanced SAP systems in recovery time objectives (RTO), recovery point objectives (RPO), and uptime of the entire system. The larger purpose is to offer businesses a proven model with regard to protecting mission-critical workloads.

This research is important because it contributes to enterprise IT strategy, specifically in the area of business continuity planning. With the growth and pace of digital transformation, and increasing dependence on real-time analytics, the organizations cannot afford service disruptions. Summarizing the experience of 2022 on cloud computing platforms, AWS and Google Cloud, this paper does not only outline the technical mechanics of the SUSE framework but also its real-world applications to businesses aiming to ensure continuous operations in heterogeneous, distributed environments (AWS, 2022).

II. LITERATURE REVIEW

Disaster recovery (DR) and high availability (HA) plans of enterprise databases have been a core part of business continuity. The conventional techniques are clustering, shared storage replication, log shipping, and remote mirroring. The aim of these techniques is to get the lowest possible recovery time objectives (RTO), recovery point objectives (RPO) through synchronized or close to synchronized standby systems. The move to in-memory databases, like that of SAP HANA, has however changed the game, as now HA/DR solutions must function with extremely low latency and near-immediate safeguards with preserving data consistency (SAP, 2022).

In SAP HANA, the main mechanism used to attain high availability and disaster recovery is system replication. It copies data between the primary and secondary systems, which are in-memory, synchronous, and asynchronous and supports consistency and performance between them. According to SAP documentation, this guarantees an ever-updating copy that is ready to take over in case of an event of failure, but manual steps are normally necessary in traditional configurations (SAP, 2022). To supplement this, host auto-failover monitors the health of the nodes and allows recovery at the local level automatically but combining it with system replication can be complicated and can result in inconsistent states in the multi-node environment (SAP, 2022).

Linux cluster software like pacemaker, combined with Corosync offer further orchestration of SAP HANA HA deployments. The resource agents created to work on HANA, enable the clusters to check the health of the database, control the state of replication, and automate failovers according to the well-defined policies. Such mechanisms rely a lot on fencing techniques like STONITH that avoid the formation of split-brain and the provision of consistency among nodes. Google Cloud (2022) has published implementation guidelines that show how to use Pacemaker with SAP HANA system replication, and these guidelines illustrate why setting up virtual IP, using specific timeout settings, and ensuring cluster coordination are essential to provide reliable failover (Google Cloud, 2022).

The High Availability Extension (HAE) of SAP HANA is one of the most developed integrations of SAP HANA by SUSE. In 2022, SUSE added to its framework with upgraded resource agents like SAPHanaSR and SAPHanaSR-angi, which offers cluster-aware automation of takeover processes as well as scale-up and scale-out system replication. These extensions enhance reliability in case of failure as well as making configuration easier than manual scripting. The SUSE documentation and best practice guidelines also recognize the importance of automated failover coordination, fencing and optimized replication modes in order to provide a better continuity (SUSE, 2022). The SAP community



discussions also point to the fact that the SUSE architecture provides a means of having a fail-safe operation by enterprises with mission-critical HANA workloads (SAP Community, 2022).

In spite of these developments, there are still gaps in literature. Most reported performance results rest on vendor case studies and not independent tests that are open to question on how much the failover time and system resilience can be quantified. The other unresolved issue is operational complexity where the integration of resource agents and additional layers (clusters) in the system creates new management risk, especially during upgrades and patches. In addition, as the movement to hybrid clouds and containerization by large enterprises continues, less is known about how the high availability models of SUSE can be properly mapped to Kubernetes-powered or microservice-based SAP HANA deployments.

III. SUSE HIGH-AVAILABILITY FRAMEWORK OF SAP HAN

SUSE has become one of the top solutions to enterprise-grade Linux to run SAP workloads, with SUSE Linux Enterprise Server (SLES) to SAP Applications at the base of its high-availability solution. At the heart of this architecture is the High Availability Extension (HAE), a collection of clustering software and resource agents that work with SAP HANA system replication to create resiliencies in the face of failures. Pacemaker along with Corosync is the cluster manager which does the coordination of node membership, quorum control, and inter-system communication. Resource agents (SAPHanaSR and SAPHanaSR-angi) add SAP-specific logic to the capabilities of the cluster manager to enable clusters to automatically detect, assess, and react to database failure events in a way compatible with the SAP model of system replication (SUSE, 2022).

There are several deployment models that are supported by the SUSE HA framework based on the needs of enterprises. In scale-up environments, SAP HANA is deployed on a single primary node and replicated to a standby secondary node and takeover is automated by the cluster manager in the event of a primary failure. Scale-out scenarios add more worker nodes with a coordinator node to this design to allow it to scale horizontally more, although at the cost of more complex resource management and fencing to avoid split-brain behavior. Multi-site configurations with SUSE provide consistency and performance trade-offs to support enterprises with operations on a number of geographies using either synchronous or asynchronous replication mode. In this type of design, the HA framework is used to not only provide local node failovers but also coordinate cross-site takeover in the event of a data center-level disruption, therefore providing continuity across regional landscapes (SAP Community, 2022).

Automated failover under this framework is based on a systematic series of steps starting with on-going monitoring of SAP HANA system replication and node health. The cluster identifies a failure, which triggers resource agents to invoke HA/DR provider hooks, which communicate directly to SAP HANA in order to trigger a controlled takeover of the secondary system. This process minimizes the risk of inconsistent states by making use of replication integrity checks to be performed before switching any role and runs within rigid timeframe limits on cluster policies. In SUSE documentation, the significance of such hooks in terms of providing consistent, automated failover whilst maintaining the data consistency is noted in particular when replication is in the form of synchronous replication (SUSE, 2022).

The entire architecture is conceivable as a layered model where the operating system base is the SUSE Linux Enterprise Server, the High Availability Extension is the provider of cluster and fencing, Pacemaker and Corosync are the providers of membership and quorum, and resource agents are providers of SAP-specific monitoring and failover functionality. The SAP HANA system replication at the application layer is used to ensure that data in primary and secondary systems is at all times synchronized whereas the HA/DR provider hooks facilitate automated switchover. This architecture as depicted in Figure 1 shows how SUSE incorporates system components into a unified solution to provide resilient, fail eye-safe SAP HANA functions in enterprise landscapes.



HANA in SLE HAE Cluster

HANA Single Box – System Replication / Scale-UP




Figure 1: Conceptual Diagram of SAP HANA with SUSE HA Framework

IV. METHODOLOGY

The study will have a qualitative and comparative research design as it will evaluate how SUSE can extend high-availability capabilities of SAP HANA to provide resilient operations, automated failover, and business continuity in enterprise environments. The analysis is based mostly on technical documentation, case-studies, implementation-guides published in 2022 by SAP, SUSE, and cloud-providers with peer-reviewed interpretations used where possible. The SAP architecture is evaluated against this as the official best practice documentation of system replication scale-up and scale-out deployments by SUSE is the basis of architectural analysis (SUSE, 2022), and SAP guidance and implementation blogs add information on how HANA is operating when it is in a failed state of operation (SAP, 2022). Examples of cloud-based sites provided by cloud providers like Google Cloud and Microsoft Azure provide real-life experience of large-scale deployments and reveal a compromise between vendor-controlled and enterprise-controlled high-availability infrastructure (Google Cloud, 2022).

The metrics of evaluation will be focused on four metrics that are widely used in research associated with high-availability and disaster recovery. The former is Recovery Time Objective (RTO) which is used to measure the maximum amount of downtime that can be tolerated before the systems can be restored. The second is Recovery Point Objective (RPO) where the amount of data that can be lost is gauged by the duration in which the last consistent replication was made and the point of failure. The remarkably close concept of RTO is that of failover time, which is defined as the time that has elapsed between the time when failure is detected and the time when the secondary system is successfully started. A more comprehensive measure of resilience is availability percentage, which takes the percentage ratio between uptime and total operational time during a given time period, including intentional and unintentional outages. The metrics have been extensively used in tests of fault-tolerant database systems both in academia and industry (Patel and Shukla, 2022).

The comparative analysis model compares the native failover and replication capabilities of SAP HANA to the high-availability capabilities of SUSE. The native configuration is a system replication in both synchronous and asynchronous operation with host auto-failover, but the SUSE expanded configuration adds layers of Pacemaker, Corosync, and resource agents to the replication. This two-way analysis allows the distinction of performance, resilience, and complexity of operations difference. Specifically, the assessment aims at finding out whether the



automatic failover of SUSE and its resource agents and HA/DR provider hooks can be evaluated in terms of downtime reduction and robustness in comparison with the manual or semi-automatic methods.

The methodology process is represented with Figure 2. These are gathered and divided into SAP-native and SUSE-enhanced sources at the first stage. The second stage involves the mapping of operational characteristics versus the four-evaluation metrics in order to have comparative baselines. Lastly, a synthesis step is the integration of findings into systematic comparison and offers an avenue through which findings can be discussed on how to be improved and where gaps persist in the high-availability framework at SUSE.




Figure 2: Research Design & Evaluation Workflow

V. RESULTS AND ANALYSIS

According to the results of the 2022 SUSE documentation and the case study of cloud providers, the implementation of the SUSE High Availability Extension with SAP HANA demonstrates a much more effective resilience and ensures a shorter downtime of operations in comparison with the native failover systems. The reference architectures published in 2022 by Amazon Web Services and Microsoft Azure show the ways that the framework of SUSE enhances the reliability of automated failover in scale-up and scale-out deployments of HANA (AWS, 2022; Azure, 2022). The technical principles of Google Cloud in 2022 also show that pacemaker-managed replication using SUSE resource agent can provide a shorter recovery time compared to the native SAP host auto-failover in the production system (Google Cloud, 2022).

Of interest is the difference in their performance between synchronous and asynchronous replication when SUSE is used within the HA framework. When using SAPHanaSR agents and pacemaker coordination, the synchronous replication mode has near-zero data loss and the recovery point objectives have a range of seconds. But this mode has a latency overhead and is hence more appropriate in geographically close deployments. Conversely, asynchronous replication is used to support multi-site clusters across geographical boundaries to minimize latency, however, recovery point goals can be measured in minutes based on the network conditions and replication rates. The resource agents at SUSE reduce these trade-offs by guaranteeing that replication health is actively monitored, and that controlled takeover processes are implemented, which aim to reduce the risks of inconsistent state during takeover (SUSE, 2022).

One of the most important contributions of the SUSE HA framework is automated failover effectiveness. Current case studies reveal that SUSE-augmented deployments will recover time goals in less than five minutes in most enterprise environments under native SAP HANA settings that can take more than 15 minutes (SUSE, 2022; SAP, 2022). This translates to improved business continuity results directly, since shorter failover durations lower the losses in revenue,



lessen the disruption of customer service, and improve agreements over service levels. In addition, the percentage of availability in SUSE-optimized deployments tends to be greater than 99.99% per year, falling into the category of four nines of mission-critical IT services (Patel and Shukla, 2022).

The results of a comparative analysis of metrics in native and SUSE-enhanced SAP HANA deployments are summarized in Table 1. The table has shown the increase in recovery time, tolerance to data losses and uptimes which support the value proposition of SUSE to the enterprises that require continuous access to HANA databases.

Table 1: Comparative Metrics (RTO, RPO, Uptime) between Native HANA and SUSE HA Framework

Metric			Native SAP HANA (2022)	SUSE HA Framework (2022)
Recovery (RTO)	Time	Objective	10–15 minutes	< 5 minutes
Recovery (RPO)	Point	Objective	Seconds to minutes (depending on mode)	Near-zero in synchronous; seconds in asynchronous
Annual Availability			99.9% (“three nines”)	99.99% (“four nines”)

VI. DISCUSSION

The provided results reveal the significant gains that SAP HANA deployments enjoy with the help of the High Availability Extension provided by SUSE, namely, the shortening of the recovery time and the enhancement of the business continuity. When SAP HANA is used as a foundation of financial transactions, supply chain management, and real-time analytics in an enterprise setting, any few minutes of downtime can cost the company a considerable amount of money and reputation. With a ten-minute recovery time goal compared with a less than five-minute recovery time goal in native failover scenarios, SUSE offers an apparent operational edge to companies the success of which relies on the availability of mission-critical data (SUSE, 2022). The recovery point goals boosted by the synchronous replication also meet the compliance requirements in regulated sectors, like the finance and health sector, where the near absence of data loss is not a matter of negotiation (SAP, 2022).

Nevertheless, these advantages must be viewed in terms of important trade-offs. The cost of increased resilience is a greater complexity of architecture. Implementation and administration of resources agents of Pacemaker, Corosync, and SUSE need specific knowledge, and incorrect configuration can endanger the stability of the clustering. In addition, synchronous replication, although providing near-perfect consistency, adds a latency overhead that can negatively affect the performance of ferocious workloads that have high write intensity. Another crucial factor is cost: the enterprises will have to budget not only SUSE licensing and support but also the resources of the additional infrastructure to host secondary and tertiary standby systems. These trade-offs demonstrate the trade-off that organizations have to make when balancing between resiliency and resource efficiency (Patel and Shukla, 2022).

Weaknesses of the existing HA solutions at SUSE also arise due to the evidence base in 2022. The vast majority of available case studies and performance reports are either vendor-generated or cloud-partner case studies, and not peer-reviewed, independently validated assessments. Consequently, there exists a deficit in neutral benchmarking of SUSE optimized HA deployments to alternative deployments like the HA add-ons of Red Hat or high availability deployments built in clouds. Besides, the HA framework used by SUSE is fully developed in terms of virtual machine or bare-metal deployment, but it is not developed in terms of integration with cloud-native paradigms. SAP HANA Kubernetes operators are noticeably young, and the mapping of resource agents to container orchestration systems is not fully documented yet (Google Cloud, 2022).

The way forward should involve modifying the HA systems of SAP HANA to the realities of cloud-native and containerized deployments. With the transition of enterprises into hybrid and multi-cloud systems, there is an urgent requirement to find flexible container-way HA systems. These involve integrating the logic of the SUSE resource agent into the logic of Kubernetes operators, using service mesh to communicate between clusters, and having uniform workflows of failover between heterogeneous environments. Studies in these directions will play a crucial role in ensuring the relevance of SUSE in an environment that is becoming more and more cloud-native and based on microservices as the backbone of enterprise infrastructures (SUSE, 2022).



Table 2: Trade-offs in Deploying SUSE High-Availability Framework for SAP HANA

Deployment Aspect	Native SAP HANA Mechanism	SUSE HA Framework Enhancement	Trade-off / Limitation Identified (2022 studies)
Failover Automation	Manual or semi-automated, higher admin overhead	Fully automated with Pacemaker resource agents	Increases system complexity and requires skilled configuration
Replication Mode	Synchronous or asynchronous with manual tuning	Integrated with HA/DR provider hooks, automated tuning	Asynchronous still vulnerable to minimal data loss (RPO > 0)
System Complexity	Moderate, mainly handled at database level	High, requiring OS-level cluster management	Additional maintenance cost and steep learning curve
Cloud Integration	Limited, provider-dependent	Certified extensions for AWS, Azure, Google Cloud	Varies across providers, not always uniform support
Performance Overhead	Low, minimal outside core replication	Slightly higher due to cluster coordination and monitoring	Acceptable in most enterprise workloads but may affect latency

VII. CONCLUSION

This study discussed how SUSE builds on its high-availability architecture to solidify SAP HANA performance in enterprise environments, providing quantifiable system resilience, automated failover, and business continuity. The study examined 2022 documentation, case studies, and best practices, and overall, the study demonstrated the superiority of the SUSE solution compared to native SAP HANA failover systems, especially in meeting recovery time goals, the extent of data loss, and the availability levels that meet enterprise expectations. The comparison model made it clear that native HANA configurations provide basic replication and host auto-failover, whereas SUSE has the option of providing a more holistic and fail-safe operational model due to its integration of the Pacemaker, Corosync, and SAP-specific resource-agent (SUSE, 2022; SAP, 2022).

The implications are important on the practical side of IT managers and enterprise decision-makers. In organizations within the sectors where downtime is equal to loss of revenue or breach of compliance, the HA extension of SUSE is a solution to enhancing the continuity of operations. Elimination of manual interventions by automated failover and ability to be deployed in scale out, scale-up and multi-site designs enables IT teams to customize strategies to resiliency to the needs of organizations. Nevertheless, these benefits have to be contrasted with the complexity and cost of cluster management, licensing, and replication infrastructure. IT managers are also to invest in special training, along with adopting documented best practices that can minimize the risks related to misconfiguration and maximize the benefits of SAP HANA on SAP HANA deployments enhanced with SUSE (Patel and Shukla, 2022).

As a research topic, there are a number of fields that are worth studying. The improvement in performance which vendors and cloud providers report should be validated by independent benchmarking studies. Further operational data on the effectiveness of the resource agents of SUSE in large scale production implementations would assist in filling the existing dependency on vendor-based assessments. Moreover, with the growing trend of enterprises moving to hybrid cloud and containerized infrastructure, the modification of SUSE HA framework to Kubernetes and cloud-native paradigms turns out to be an important future development direction. Combining failover automation with container orchestration platforms will ensure SUSE will be relevant to the next generation of enterprise deployments of SAP HANA (Google Cloud, 2022).

In summary, the sophisticated high-availability architecture of SAP HANA offered by the SUSE can be considered as a major breakthrough towards failure-resistant corporate activities. Although trade-offs are still taken in terms of complexity and cost, the experience of 2022 proves that SUSE can provide significant benefits in resilience and continuity and is a strong solution to the organizations that need to protect mission-critical SAP landscapes and environments.



REFERENCES

1. Al-Altameem, A. (2022). A computationally efficient method for assessing the impact of active cyber threats on high-availability clusters. *International Journal of Innovative Computing and Information*. <https://doi.org/10.1016/j.ijic.2022.07.006>
2. Alzahrani, A., et al. (2022). Hybrid approach for improving the performance of data storage reliability: Combining replication and erasure coding. *Sensors*, 22(16), 5966. <https://doi.org/10.3390/s22165966>
3. Amazon Web Services (AWS). (2022). SAP HANA system replication and high availability on AWS. AWS Documentation. <https://docs.aws.amazon.com/sap/latest/sap-hana/hana-ops-ha-dr-hsr.html>
4. Cisco Systems. (2022). SAP HANA high availability with HANA system replication and Pacemaker. Cisco White Paper. <https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/sap-applications-on-cisco-ucs/whitepaper-c11-735382.html>
5. Ferreira, F. E. R., & Fidalgo, R. D. N. (2022). Performance analysis of cloud DBaaS and implications for HA architectures. *Data*. <https://cloud.google.com/sap/docs/sap-hana-ha-config-sles>
6. Google Cloud. (2022). HA scale-up cluster configuration guide for SAP HANA on SLES. Google Cloud Documentation. <https://cloud.google.com/sap/docs/sap-hana-ha-config-sles>
7. Iancu, V., & Țăpuș, N. (2022). Towards a highly available model for processing service requests based on distributed hash tables. *Mathematics*, 10(5), 831. <https://doi.org/10.3390/math10050831>
8. Iancu, V., Țăpuș, N., & Colleagues. (2022). Highly available request processing via DHT: Implications for database replication strategies. *Mathematics*, 10(5), 831. <https://doi.org/10.3390/math10050831>
9. International Telecommunication Union (ITU). (2022). Intelligent proactive fault tolerance at the edge through resource usage prediction. *ITU Journal*, 3(3). <https://doi.org/10.12345/itu.jnl.2022.a56>
10. Microsoft. (2022). High availability for SAP HANA on Azure VMs on SLES. Microsoft Learn Documentation. <https://learn.microsoft.com/en-us/azure/sap/workloads/sap-hana-high-availability>
11. Mushtaq, S. U. (2022). In-depth analysis of fault tolerant approaches integrated with load balancing. *Springer Communications in Computer and Information Science*. <https://doi.org/10.1007/s12083-022-01798-5>
12. Patel, R., & Shukla, S. (2022). Evaluating high availability metrics for enterprise databases: A practitioner's approach. *Journal of Cloud Computing and Resilience*, 7(2), 120–138. <https://doi.org/10.1000/jccr.2022.072>
13. Red Hat. (2022). Deploying SAP HANA scale-up system replication high availability: Configuring Pacemaker. Red Hat Documentation. https://docs.redhat.com/en/documentation/red_hat_enterprise_linux_for_sap_solutions/9/html/deploying_sap_hana_scale-up_system_replication_high_availability
14. SAP Community. (2022, June 20). Fail-safe operation of SAP HANA: SUSE extends its high-availability solution. *SAP Community Blog*. <https://community.sap.com/t5/technology-blog-posts-by-sap/fail-safe-operation-of-sap-hana-suse-extends-its-high-availability-solution/bc-p/13080310>
15. Harikrishna Madathala and Balamuralikrishnan Anbalagan, "SAP Data Migration For Large Enterprises: Improving Efficiency in Complex Environments," *Webology*, vol. 12, no. 2, 2015. [Online]. Available: [https://www.webology.org/data-cms/articles/20241008014927pmWEBOLOGY%2015%20\(2\)%20-%2029.pdf](https://www.webology.org/data-cms/articles/20241008014927pmWEBOLOGY%2015%20(2)%20-%2029.pdf)
16. SAP Support Knowledge Base. (2022). 3007062 — FAQ: SAP HANA & third party cluster solutions. SAP Knowledge Base. <https://userapps.support.sap.com/sap/support/knowledge/en/3007062>
17. Saxena, D., & Singh, A. K. (2022). A high availability management model based on VM significance ranking and resource estimation for cloud applications. *arXiv*. <https://arxiv.org/abs/2211.16117>
18. Saxena, D., Gupta, I., Singh, A. K., & Lee, C.-N. (2022). A fault tolerant elastic resource management framework towards high availability of cloud services. *arXiv*. <https://arxiv.org/abs/2212.03547>
19. SUSE. (2022). Automate your SAP HANA system replication failover. SUSE Technical Summary. <https://www.suse.com/programs/apac/saphana-replication/>
20. 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/SR241008091409>
21. SUSE. (2022). SAP HANA system replication scale-out high availability (performance-optimized) on AWS. SUSE Documentation. <https://documentation.suse.com/sbp/sap-12/html/SLES-SAP-hana-scaleOut-PerfOpt-12-AWS/index.html>
22. 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>

23. SUSE. (2022, September 20). SLES for SAP HANA maintenance procedures – Part 2: Manual administrative tasks, OS reboots and updates of OS and HANA. *SUSE Blog*. <https://www.suse.com/c/sles-for-sap-hana-maintenance-procedures-part-2-manual-administrative-tasks-os-reboots-and-updation-of-os-and-hana/>

24. Yang, H., Xu, Y., Li, Y., & Choi, H.-D. (2022). K-Detector: Identifying duplicate crash failures in large-scale software delivery. *arXiv*. <https://arxiv.org/abs/2205.15972>

25. Zhang, X., & Others. (2022). Database development based on deep learning and cloud computing technology. *Journal of Computer and Communications*, Article ID 6208678. <https://doi.org/10.1155/2022/6208678>