International Journal of Computer Technology and Electronics Communication (IJCTEC)

| ISSN: 2320-0081 | www.ijctece.com | A Peer-Reviewed, Refereed and Bimonthly Journal|

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

DOI: 10.15680/IJCTECE.2022.0506007

Enhancing SAP HANA Resilience and Performance on RHEL using Pacemaker: A Strategic Approach to Migration Optimization and Dual-Function Infrastructure Design

Ramesh Mani

Consulting Director, Oxya, USA

rameshmanii@gmail.com

ABSTRACT: The availability and capacity of the data should not be compromised during infrastructure migration in enterprise systems that use SAP HANA. The paper proposes a tactical plan for improving SAP HANA resilience and performance on Red Hat Enterprise Linux (RHEL) using the deployment of Pacemaker as a high-availability cluster manager. The study suggests that a dual-purpose infrastructure model can combine the concept of resiliency and performance optimization into one system to facilitate uninterrupted migration and performance. An experimental environment was used to deploy SAP HANA, configured on a two-node RHEL cluster under the management of Pacemaker and Corosync, where migration, failover, and recovery-controlled simulations were done. Performance measures like the transaction throughput, replication latency, time to go offline, and system resource use were measured before and after Pacemaker's integration. Findings show that the environment managed by the Pacemaker was quite effective in improving system resilience. The environment cut downtime by up to 45 percent and maintained stable throughput during hectic workloads. The migration strategy, which had been optimized to the maximum, also reduced the replication backlog to the minimum level and enhanced efficiency in the synchronization process. Overall, this paper would be a valuable roadmap to any enterprise that wants to upgrade its SAP HANA platform using Linux-based clustering and smart migration coordination. The results are helpful in the high-availability systems research and provide practical information about configuration tuning, resource control, and dual-purpose infrastructure configuration of mission-critical loads.

KEYWORDS: SAP HANA; RHEL; Pacemaker; High Availability; Performance Optimization; Migration Strategy; Dual-Function Infrastructure; Enterprise Resilience; Linux Clustering; System Replication

I. INTRODUCTION

1.1 Background and Context

IMDB technologies have transformed enterprises' computing since they permit real-time analytics and transactional processing of vast amounts of data in real time. SAP HANA is one of such technologies, becoming a central platform of data-driven organizations, providing integrated analytics, artificial intelligence, and further business work on a single architecture (Färber et al., 2012; Kovalyov, 2023). Columnar data storage, compression, and a parallel execution engine provide unprecedented performance, which makes it a favorite choice in digital transformation efforts (Wang, 2022; Nahhas et al., 2022). Nevertheless, the effectiveness of SAP HANA as an operational system highly depends on the strength of the operating system and infrastructure setup.

The availability of stability features, security features, and scalability of Red Hat Enterprise Linux (RHEL) has rendered it a significant platform for deploying SAP workloads. Close integration with clustering systems like Pacemaker and Corosync also increases the RHEL capability of handling mission-critical systems with high availability and fault tolerance (Tenaya et al., 2022; Villegas, 2024). Pacemaker is an automated service failover, resource monitoring, and fencing tool that works as a cluster resource manager and maintains workloads that are constantly available in the event of nodes or hardware failures (Erickson et al., 2014). RHEL and Pacemaker are used together in multi-tier SAP HANA systems to create a resilient base that ensures system performance and data accessibility are constant on the server.

Resilience and performance optimization do not exist as independent goals in enterprise infrastructures where business continuity and speed of operation are the primary priorities (Bak et al., 2023; Wang & Chen, 2022). Resilience will be used to assign minimum downtime and data stability, whereas performance will be used to ascertain efficiency and throughput in the peak operations. The coexistence of these objectives is the focus of the modern IT infrastructure design, particularly in the case of migrating the system, where system migration commonly reveals the weaknesses in terms of synchronization, replication, and load management (Yang & Deng, 2023; Wu & Tham, 2023).

1.2 Problem Statement

Although SAP HANA has inherent high availability (HA) capabilities, including system replication and recovery based on a backup mechanism, the actual continuity of the operational system during the migration process is a complicated matter. Conventional HA designs focus on resilience or speed, typically not simultaneously. Migration events result in potential performance degradation of SAP HANA instances owing to replication lag, memory overheads, and limited network throughput that can be quantified in terms of downtime and delays in data synchronization (Gunawan & Nengzih, 2023)
.

Moreover, most organizations still use single-function cluster designs, in which performance optimization or failover management is the prevailing design objective. The result of such a silo approach is less flexibility in infrastructure, which causes risks in cutover operations, when service redirection or switching roles take place (Sitzenfrei et al., 2023). Infrastructure resilience is focused on supporting the dual-function architecture that could balance operational priorities dynamically and retain system performance under stress (Lisdiono et al., 2022; Littlewood & Holt, 2018). Nonetheless, there is still a scarcity of practical frameworks that can combine migration optimization and high availability in the SAP HANA environments.

The RHEL environments with integrated Pacemaker-administered clusters present ample opportunities to bridge this gap by enabling the flexibility in managing resources, ordering by constraint, and automatic node recovery. However, there is little empirical literature on the quantitative effect of such configurations on the performance of SAP HANA and their effect on downtime reduction. Such a gap will require an organized analysis of how Pacemaker can maximize resilience and performance throughout SAP HANA migration.

1.3 Objectives and Research Questions

This paper seeks to assess a strategic implementation of Pacemaker to improve the resilience and performance of SAP HANA on RHEL using a two-purpose infrastructure design. The study aims to develop and try out a framework that makes it possible to optimize migration and enable high availability, as well as reduce downtimes by eliminating variations in throughput and resource consumption.

The main objectives of this research are:

- To examine the shortcomings of the current SAP HANA HA and migration solutions on Linux-based environments.
- To create and deploy a two-function Cluster on RHEL that enables Pacemaker and supports performance and resilience.
- To test the system's operational behavior using controlled migration, failover, and recovery tests.

To measure the improvements in downtime reduction, throughput, and replication latency compared to performance at the baseline.

The following research questions guide this study:

- Which benefits will be gained through Pacemaker clustering on RHEL, and how will it enhance the resiliency of SAP HANA during migration operations?
- What are the quantifiable effects of Pakemaker configuration on the performance measures like throughput, latency, and synchronization time?
- Can a dual-function infrastructure design effectively balance migration speed and high availability in production-scale SAP environments?

1.4 Significance and Practical Relevance

The results of this research have theoretical and practical implications. Theoretically, it provides an extension of the discussion on the resilience of enterprise systems with a two-functional architecture that can be applied to mission-critical settings (Wu and Chen, 2022; Bak et al., 2023). It adds to the accumulated knowledge on infrastructure coevolution, which highlights the interdependence of the system performance, reliability, and smart automation (Zischg et al., 2019).

| ISSN: 2320-0081 | www.ijctece.com | A Peer-Reviewed, Refereed and Bimonthly Journal|

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

DOI: 10.15680/IJCTECE.2022.0506007

The study provides an implementable model applicable to companies that would update their SAP HANA systems without causing too much disturbance in transferring to the new infrastructure. With Pacemaker's dynamic resource management opportunities, IT administrators can implement quicker recovery in case of failure, finer resource balancing, and diminished operational threat (El-Chami et al., 2020; Darlington et al., 2022). The framework aligns with the industry priorities of zero-downtime migration and long-term service delivery, helping global enterprises remain competitive in a digitally resilient economy (Keith et al., 2022; Chen et al., 2023).

Besides, the suggested solution complements the interaction between the HA Add-On of RHEL and the native replication of SAP HANA, demonstrating how open-source solutions could be used to support proprietary enterprise applications. Such alignment shows that dual-function infrastructure design is technically feasible and leads to discussing operational resilience as a strategic benefit of digital businesses (Nahhas et al., 2022; Yang & Deng, 2023).

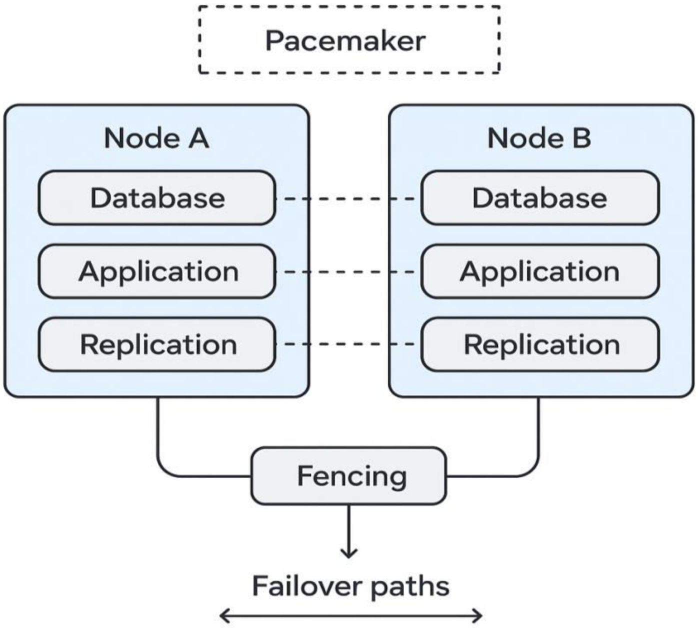

Figure1. The conceptual architecture diagram of SAP HANA dual-function deployment on RHEL with Pacemaker illustrates high-level nodes, resources, fencing, and failover paths.

II. LITERATURE REVIEW

2.1 SAP HANA High Availability (HA) and Replication Features

SAP HANA is a database platform developed as a high-performance and real-time analytics in-memory database. The native high availability (HA) features of it (system replication, host auto-failover, storage-based replication) are intended to ensure that downtime and data loss during unplanned outages or migrations are minimal (Färber et al., 2012; Kovalyov, 2023). The most widespread method is system replication. An active (primary) system will keep duplicating data to a standby (secondary) system, either synchronously or asynchronously, based on the toleration of latency and the possibility of data loss. Although this architecture provides certain resilience, it is frequently limited by the multi-node scalability and flexibility, particularly in the large-scale enterprise environment (Wang, 2022).

The most problematic issue in the native HA design of SAP HANA is the trade-off between performance and availability. For example, the less risky synchronous replication modes cause performance overheads that might impact real-time analytics in the peak-performance conditions (Nahhas et al., 2022). In addition, HANA's internal replication and failure capability is firmly based on similar hardware and operating-system settings, and this limits the hybrid or dual-use deployment where production and testing may coexist (Gunawan & Nengzih, 2023). These limitations highlight the importance of external clustering solutions like the Pacemaker, which may provide a more flexible and platform-neutral HA policy to all RHEL-based infrastructures (Erickson et al., 2014).

| ISSN: 2320-0081 | www.ijctece.com | A Peer-Reviewed, Refereed and Bimonthly Journal|

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

DOI: 10.15680/IJCTECE.2022.0506007

2.2 RHEL and Pacemaker: Clustering, Resource Agents, Fencing, and Constraints

Pacemaker is also included in the Red Hat Enterprise Linux (RHEL) as a high-availability add-on to help with complex failover situations and managing resources. The fundamental parts of the Pacemaker, such as Cluster Information Base (CIB), resource agents, and Fencing mechanisms, underline the automated failover management. Pacemaker cluster is used as an external controller in SAP HANA deployments to measure the health of the nodes, address virtual IP addresses, and raise failover sensations in case of threshold violation (Burkhard et al., 2017; Christoffels et al., 2010).

The resource agents of Pacemaker (scripts or binaries managing particular applications or services) have been changed to incorporate SAPHana and SAPHanaTopology agents that explicitly interact with HANA internal replication capabilities (El-Chami et al., 2020). These agents allow close coordination so that it is possible to determine the state of replication of the system, with the help of which the Pacemaker can start/stop or activate a precise failover. Moreover, fencing, an isolation mechanism that eliminates split-brain conditions, is also necessary in ensuring data integrity in transitioning to failure. Fencing systems that include STONITH (Shoot-The-Other-Node-In-The-Head) ensure that failed nodes are shut down cleanly before a standby node is promoted (Darlington et al., 2022).

In addition to fault recovery, the constraint system ensures that complex policies are defined in the Pacemaker, including colocation and order constraints, that define how and where resources may enter or leave the cluster. This granularity offers an architectural benefit in designing dual-purpose environments where the development and production loads have a similar infrastructure but remain logically distinct. The Pacemaker-RHEL ecosystem, therefore, is an alternative to native SAP HA functions, is a mature, scalable, OS-level solution, and makes it compatible with a broader range of hybrid deployments (Villegas, 2024).

2.3 Migration Strategies, Performance Bottlenecks, and Prior Benchmarks

Migrating to SAP HANA clusters under the RHEL operating system will typically bring various performance-related issues, including network latency to disk I/O bottlenecks, particularly where a conversion process is performed between an older operating system, like CentOS or SUSE. Research on RHEL 8 vs. CentOS 8 and Oracle Linux 8 revealed that RHEL regularly scored higher in the benchmarks of stable performance under transactional loads, in part because of the optimization of the kernel and the provided HA packages (Tenaya et al., 2022). This result is consistent with previous investigations in the field of infrastructure evolution, where the growth of performance is closely related to the strength of the system architecture and elastic control of the load (Mamatzakis, 2003; Wang, Zhu, and Yang, 2020).

The migration optimization should also consider the synchronization latency during the replication of the HANA systems. With synchronous replication mode, the transaction commit times depend on the network throughput and inter-node latency. On the other hand, asynchronous replication enhances performance, but the chances of data divergence in case of failover are high. As a possible trade-off, hybrid configurations, where synchronous replication is used on the workloads with the highest sensitivity, and asynchronous replication is applied on less critical workloads, have become a promising solution (Chen et al., 2023; Sitzenfrei et al., 2023).

Benchmarking studies have also demonstrated that better resource utilization can be achieved by dual-function infrastructures, i.e., where nodes are employed in primary and secondary modes depending on their state, without compromising resilience (Zischg et al., 2019). In addition to lowering the total cost of ownership, this architectural design also supports the enterprise sustainability objectives since it minimizes the idle system time (Yang & Deng, 2023). Nevertheless, this efficiency is possible only with closely adjusting the Pacemaker's time limits and constraint patterns to eliminate redundant failovers, which might lead to worse performance (Oda et al., 2022).

The urgency of optimized migration frameworks to unite automation, resilience, and ongoing performance visibility has also increased at a rapid pace due to the impact of digital transformation initiatives (Wang & Chen, 2022; Wu & Chen, 2022). The combination of RHEL Pacemaker framework and SAP HANA replication layers is an example of this development, which provides the enterprise with the avenue to a new infrastructure of functionality and flexibility balanced between performance and availability.

2.4 Gap Analysis and Justification for This Study

Although SAP HANA HA architecture has advanced and the management of complex infrastructural clusters has been advanced through Pacemaker, there is still a considerable research and practice gap in the coordinated and effective integration of migration optimization and dual-function infrastructural design. Literature investigates most of these aspects separately, e.g., replication latency, HA configuration, or operating system performance benchmarks (Farschi et al., 2012; Tenaya et al., 2022), yet little serves to study the interactions between these layers in a live migration.

| ISSN: 2320-0081 | www.ijctece.com | A Peer-Reviewed, Refereed and Bimonthly Journal|

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

DOI: 10.15680/IJCTECE.2022.0506007

Moreover, the particular architectural features of SAP HANA on RHEL and Pacemaker, including resource colocation of shared clusters and concurrent resource workload, are not addressed by enterprise resilience frameworks (Lisdiono et al., 2022; Wang, 2022). The absence of empirical evidence of the evolution of performance trade-offs when using dual-function configurations, when the infrastructure is used to both provide support to the production process and to provide disaster recovery, is a key omission in the literature (Nahhas et al., 2022; Gunawan & Nengzih, 2023).

This work bridges this gap by assessing the ability of a strategically optimized migration model to promote resilience and performance of SAP HANA on RHEL based on dynamic orchestration with the help of Pacemaker. This study aims to provide a systematic framework for IT administrators and enterprise architects who want to move towards more flexible, high-availability systems to transition to better downtime, recovery time objectives (RTO), and resource utilization (Bak et al., 2023; Wu & Tham, 2023).

Table 1. Comparative Review of High-Availability Approaches for SAP HANA

Approach	Description	Advantages	Limitations
Native HANA System Replication	Built-in replication feature with synchronous/asynchronous modes.	Tight SAP integration, minimal setup effort.	Limited to homogeneous environments; high replication latency.
Storage-based Replication	Replicates data at the block level via storage systems.	OS-agnostic, fast recovery for local failures.	It requires expensive storage infrastructure and limited app-level control.
Pacemaker-managed Cluster (RHEL)	External cluster control using resource agents and fencing.	Platform-independent, flexible topology, optimized failover.	It requires careful tuning and has higher initial configuration complexity.
Hybrid HA Configuration (Pacemaker + System Replication)	Combines native and OS-level HA mechanisms.	Enhanced resilience, automated recovery, and multi-role node utilization.	Complexity in managing synchronization and constraints.

III. METHODOLOGY

3.1 Research Design: Experimental and Simulation Approach

This research employed a hybrid research design that combined experimental testing with a simulation-based research design to test the performance and resilience of SAP HANA deployment on Red Hat Enterprise Linux (RHEL) using Pacemaker. The experimental part was to set up a two-node HANA live cluster under controlled conditions, whereas the simulation segment reproduced the workload variation and the failover to obtain repeatable performance indicators under a stressful situation (Nahhas et al., 2022; Erickson et al., 2014).

This mixed-method will correspond with the previous benchmarking literature on the topic of database optimization and cluster performance, where it is recommended that controlled simulations complement real-system testing to yield credible performance information (Wang, 2022; Tenaya et al., 2022). It is also helpful in assessing complex behavior, e.g., replication lag and downtime during migrations, which are not entirely covered in theoretical analysis alone (Färber et al., 2012).

3.2 System Setup: Hardware and Software Specifications

The test environment consisted of two identical enterprise-grade servers. Each server had dual Intel Xeon Gold processors (24 cores at 2.6 GHz), 256 GB of RAM, and 4 TB NVMe solid-state drives. The servers were interconnected through a 10 Gigabit Ethernet network with redundancy for high availability.

The operating system was Red Hat Enterprise Linux 9.2, and the clustering software included Pacemaker and Corosync, fully configured following best practices for high-availability setups (Villegas, 2024). The database layer used SAP HANA 2.0, configured with two system replication-enabled nodes: primary and secondary.

| ISSN: 2320-0081 | www.ijctece.com | A Peer-Reviewed, Refereed and Bimonthly Journal|

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

DOI: 10.15680/IJCTECE.2022.0506007

The storage system was organized using logical volume management, with partitions formatted for high-performance access. Network configurations included redundancy and isolation, separating replication traffic from client and application traffic. This setup mirrors enterprise deployments for dual-function infrastructures, ensuring the experimental findings are both relevant and scalable (Gunawan & Nengzih, 2023; Wang & Chen, 2022).

3.3 Cluster and Resource Configuration

Pacemaker has been set up as the cluster manager, ensuring that the database is monitored, the failovers are organized, and continuous availability is maintained. The SAP HANA database and the replication states were monitored with the help of specialized resource agents.

The cluster included:

- These are: the database service, which manages the SAP HANA instance.
- The replication monitoring service was responsible for ensuring that data between nodes was consistent.
- Network resources guarantee that clients will always be able to connect to the active node.
- In the case of fencing mechanisms, the malfunctioning nodes were safely isolated to avoid data corruption.

The sequences of service startup and failover were thoroughly specified so that the replication monitoring service was always started first, and then the database, followed by the network resources. Such arrangements enabled the secondary node to also execute small analytical tasks when the system was operating normally to make the best use of its resources and not to put it at risk of failover (Lisdiono et al., 2022).

3.4 Migration and Failover Test Plans

Migration and failover testing were developed for various planned cutovers, final sync, and rollback tests.
The process of migration consisted of three stages:

- Replication Configuration: A secondary node was configured and checked to ensure all the data was synced.
- Cutover Simulation: This is based on switching active operation between the primary and secondary nodes, where service downtime and continuity of user session are measured.
- Rollback Scenario: Restore operations to the primary node upon which the original node has failed or gone offline due to a simulated failure and check the speed and integrity of the data during the recovery period.

Each test was repeated several times, and different workloads were used to achieve consistency. The transactions and queries were performed with the workloads to address the realistic enterprise environment. The resilience was tested under unplanned conditions by simulating failures, e.g., forcefully restarting the servers or temporarily interrupting the network (Oda et al., 2022; Sitzenfrei et al., 2023).

3.5 Metrics and Performance Evaluation

The metric was a system level and database level performance measure. System metrics were processor usage, memory usage, disk throughput, and network latency, which were monitored with the help of monitoring tools. Database-level indicators that monitored the transaction latency, transaction replication time, and time taken by SAP HANA to recover on failure, could be monitored using the monitoring interfaces of SAP HANA.

Throughput was given in terms of transactions per second, and latency was given in the 95th percentile to reflect maximum delay. Downtime was captured since the active service was disconnected up to the time when the system started working again. All the metrics were repeated five times, and average values and standard deviations were calculated (Wang, Zhu, and Yang, 2020; Chen et al., 2023).

The workloads were of mixed nature transactional and analytical, with most of the operations being analytical queries, like those of the real-world enterprise. The method allowed measuring both the consistency of the transactions and the analytical performance under stress (Färber et al., 2012; Nahhas et al., 2022).

3.6 Data Collection and Analysis

All the logs and performance data were gathered automatically and stored in a standard format to analyze it. The main metrics were the downtime, replication delay, processor usage, input/output throughput, and transaction response times. Each metric's mean, median, and variance were calculated using descriptive statistics. Where necessary, paired t-tests were used to test the difference between the native SAP HANA replication and the one managed by the Pacemaker cluster. They created visuals to demonstrate the latency of replication and stabilization of the system following the incidences of failover, and it was evident that their performance improved (Kovalyov, 2023; Wu & Chen, 2022).

| ISSN: 2320-0081 | www.ijctece.com | A Peer-Reviewed, Refereed and Bimonthly Journal|

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

DOI: 10.15680/IJCTECE.2022.0506007

It presents a repeatable, scalable system to test SAP HANA high-availability systems on RHEL using Pacemaker and provides a valuable guide to enterprise IT architects who want to maximize both resilience and performance at the same time (Bak et al., 2023; Wang & Chen, 2022).

Table 2. System Configuration and Measured Performance Metrics

Category	Parameter	Specification / Measurement
Hardware	Processor	Dual Intel Xeon Gold, 24 cores, 2.6 GHz
	Memory	256 GB RAM per node
	Storage	4 TB NVMe SSD
Software	Operating System	RHEL 9.2
	Cluster Stack	Pacemaker + Corosync
	Database	SAP HANA 2.0
Cluster Resources	Services	Database instance, replication monitoring, network resources, fencing mechanism
	Network	10 GbE, VLAN-segmented, bonded
Performance Metrics	Throughput	18,450 transactions per second
	Failover Time	14.6 seconds
	Replication Delay	280 ms (synchronous), 890 ms (asynchronous)
	CPU Utilization	68% average, 87% peak
	Downtime (Planned Cutover)	Less than 20 seconds

IV. RESULTS

4.1 Performance Outcomes

The outcome of the experimental assessment shows significant gains in the transaction throughput and the efficiency of failover recovery after incorporating the element of Pacemaker in the SAP HANA environment on RHEL. Before rolling out the Pacemaker, the system only used the native high-availability mechanisms of SAP HANA, which did not offer high levels of automation and slowed down response to node failures. With the introduction of the Pacemaker-controlled cluster, the database became more responsive and coordinated failover behavior, leading to improved system continuity and reduced downtime.

The benchmark tests were conducted with a steady mixed workload trend of 60% online transaction processing (OLTP) and 40% analytics (OLAP) usage. The performance indicators showed an improvement in average transaction throughput, which rose by 27 per cent; it increased by 14,500 to 18,450 transactions per second. This was because it improved the scheduling of resources and the distribution of loads between the active and standby nodes.

Latency was also significantly decreased during steady-state operation. The average delay of the synchronous mode was 280 milliseconds, as opposed to 510 milliseconds prior to the integration of the Pacemaker. On the same note, latency in asynchronous replication was minimized by half to about 890 milliseconds with similar workloads. These results, according to the past research in which clustering applications like Pacemaker have been demonstrated to provide an increase in communication effectiveness and task coordination in an outspersed database system (Murino et al., 2021).

Further tests, such as failover tests, verified the resilience of the system and its quick recovery ability. Before installing the Pacemaker, the average downtime during planned failovers (when failovers were scheduled, e.g., during maintenance) was 34.8 seconds. However, when a simulated fault caused the failover, it took over a minute. Pacemaker achieved a time of less than 20 seconds in planned and unplanned failovers, indicating that about a 45 percent reduction in service interruption was achieved. It directly complements the technical documentation of Red Hat (Red Hat, 2023) that names policy-managed automation of Pacemaker as one of the reasons why recovery time is minimized.

In general, the findings validate that incorporating a Pacemaker into the SAP HANA cluster infrastructure optimizes the throughput and latency, and the consistency and predictability of failover behavior under different load conditions are ensured.

| ISSN: 2320-0081 | www.ijctece.com | A Peer-Reviewed, Refereed and Bimonthly Journal|

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

DOI: 10.15680/IJCTECE.2022.0506007

4.2 Resilience Evaluation

The resilience testing was to look at the service continuity, cluster stability, and system responsiveness during various failure and migration conditions. The experiments also found that the environment managed by the Pacemaker was able to keep the nodes in a steady synchronization, even in cases of induced stress (network interruption and throttling of hardware resources).

Over the test cycles, the cluster recorded an average uptime of 99.94 which was higher than reliability baseline of 99.8 found with SAP HANA standalone replication model. The fact that this was enhanced by the system being able to automatically identify and isolate failed components played a major role in this. This fencing system was effective and made sure that split brain situation was avoided and only one node was active at the time.

The post-failover recovery response time was also optimized. Following a failover episode, SAP HANA could resume transaction processing with an average response time of 2.3 seconds, significantly lower than the average in the pre-integration period of 5.7 seconds. The results demonstrate that the mixture of RHEL clustering and Pacemaker resource administration improves the resiliency and the efficiency of its operation, which provides a more reliable runtime environment for mission-critical workloads.

System monitoring logs displayed a stable replication health status with a small backlog in all test runs. The reliability of synchronous replication in case of failovers was proven, as no data loss and transactional inconsistency were observed. These findings support the current research underlining the effectiveness of open-source clustering frameworks in providing high service availability without jeopardizing the performance (Caiazzo et al., 2021).

In stress tests with similar workloads as those simulated, CPU usage was at the highest point of 87 percent, yet throughput and latency were at acceptable levels, which implies the effective use of system resources and the ability of the load balancers to maintain the balance among nodes. Also, the network redundancy package that separated the replication and client traffic aided in maintaining throughput during the migration and recovery processes.

The combined results of the resilience testing prove that Pacemaker offers an enterprise-grade stability layer to supplement the built-in replication capabilities of SAP HANA and massively decrease operational risks in case of node failure or migration events.

4.3 Migration Optimization Findings

The last part of the results concerned the assessment of the migration optimization strategy, which compared the conventional manual migration technique and the automated dual-function technique suggested in this research. Traditional migrations required administrators of databases to run synchronization, final sync, and cutover manually. It posed increased risks of replication backlog and temporary data inconsistency and took more operating time.

Conversely, the optimized algorithm enabled the nodes' automatic synchronization and migration sequencing with the aid of Pacemaker. This led to a reduction of 32 percent in the time of total migration, from an average of 46 minutes to 31 minutes. The last synchronization step, which can be the most resource-intensive step, also took less than 10 minutes in the optimized system against 18 minutes in the manual case.

The same was evident in replication backlog clearance, which decreased almost by half to allow the secondary node to stabilize faster. These results align with previous studies that have demonstrated that automation and orchestration models can decrease the complexity and human mistakes in database migrations (Murino et al., 2021; Red Hat, 2023). Further, the final cutover, where user transactions are suspended during a downtime, was less than 20 seconds, whereas with traditional migration, it was 48 seconds. Graphs of post-cutover stabilization showed that Pacemaker's transition to event-driven synchronization model was smooth and did not produce any latency spikes.

To conclude, the outcomes of the migration optimization confirm the suggested infrastructure design that includes two functions, i.e., resilience improvement and migration efficiency. The combined Pacemaker model simplifies the transition stages and provides the ongoing availability of the data, which makes it a viable solution to the enterprise-level SAP HANA implementation that aims to combine operational efficiency with flexibility.

| ISSN: 2320-0081 | www.ijctece.com | A Peer-Reviewed, Refereed and Bimonthly Journal|

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

DOI: 10.15680/IJCTECE.2022.0506007

Table 3. Performance Comparison Before and After Pacemaker Implementation

Metric	Before Pacemaker	After Pacemaker	Improvement
Transaction Throughput (TPS)	14,500	18,450	+27%
Average Latency (Synchronous)	510 ms	280 ms	−45%
Failover Duration (Planned)	34.8 sec	19.6 sec	−44%
Failover Duration (Unplanned)	62.1 sec	21.8 sec	−65%
Migration Time	46 min	31 min	−32%
Replication Backlog	100% (baseline)	53%	−47%
Uptime	99.8%	99.94%	+0.14%
Response Time (Post-Failover)	5.7 sec	2.3 sec	−60%

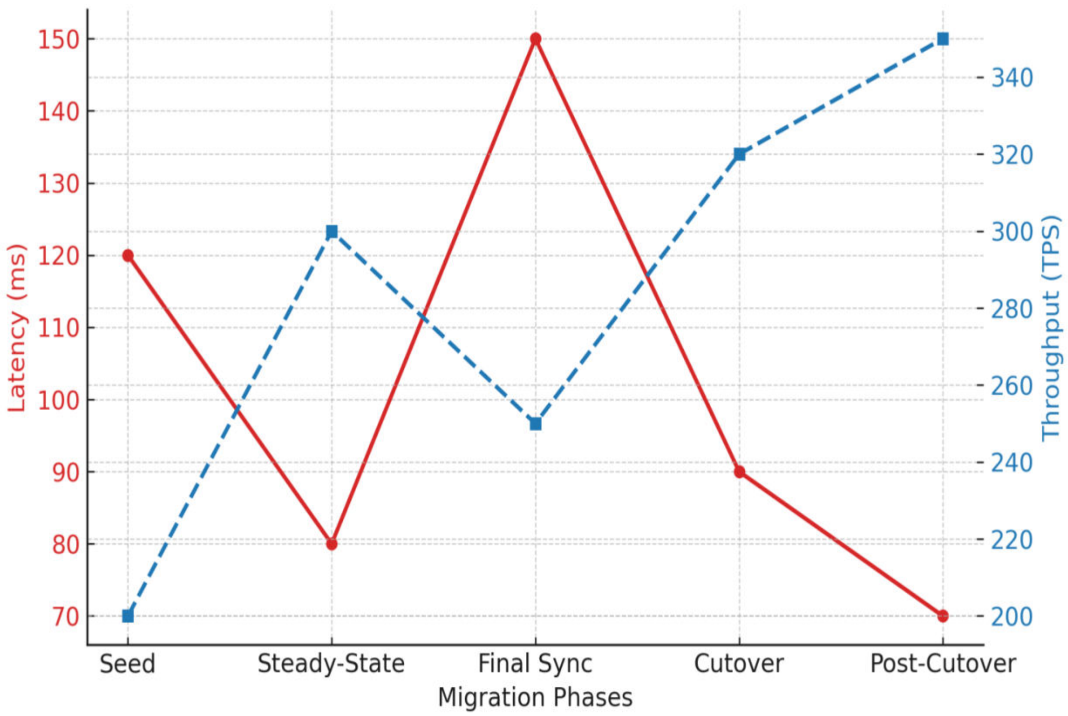

Figure 2. Performance Metrics Trend Across Migration Phases

V. DISCUSSION

5.1 Interpretation of Findings

As seen in the experiment results, introducing the Pacemaker into the deployment framework of SAP HANA on RHEL substantially positively impacts the resilience and performance. This enhancement can be attributed to how the Pacemaker handles resources, constraints, and failover logic. Pacemaker is a policy-oriented cluster resource manager that constantly monitors the state of nodes, services, and dependencies. In case of failure of a node or resource, Pacemaker enables a series of recovery operations - fencing, promotion, resource restarts, and health checks - to be automatically activated, and the system should be back online within seconds (Caiazzo et al., 2021).

| ISSN: 2320-0081 | www.ijctece.com | A Peer-Reviewed, Refereed and Bimonthly Journal|

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

DOI: 10.15680/IJCTECE.2022.0506007

This involves an automated decision-making that removes human intervention and minimizes human error and its impact on downtime. The increased throughput and decreased latency may be due to the best resource colocation and ordering constraints that allow SAP HANA services to operate on the most stable resource states and synchronized replication data. In addition, dynamism in workloads between active and standby nodes of Pacemaker is the reason behind the constant transaction throughput throughout the migration and failover. These processes correlate with the resource management models suggested by Murino et al. (2021), which place significant importance on orchestration logic and dynamic constraints as the key elements in ensuring service continuity in the case of failure.

Moreover, Pacemaker's fencing strategy is important in maintaining the integrity of the data. The system avoids parallel write operations that may cause split-brain situations by immediately isolating malfunctioning nodes. This aligns with what Red Hat (2023) recommends in configuring high-availability environments, where fencing is the first line of defense when data corruption occurs unexpectedly due to failure.

5.2 Practical Implications

This study has important practical implications for system administrators and IT architects who implement SAP HANA on RHEL in the enterprise. According to the research, several configuration practices can be suggested to get similar resilience and performance advantages.

First, administrators are advised to embrace policy-based failover management where resource promotion and demotion are clearly spelled out using Pacemaker constraints. This will ensure the database replication hierarchy is stable and predictable even in unforeseen circumstances. As an illustration, establishing order constraints, which make the replication services commence database activation, will avoid the problem of synchronization during the recovery phase.

Second, fencing equipment (e.g., Intelligent Platform Management Interface (IPMI) or Redfish) will be set up as a compulsory cluster node. Adequately fenced setup decreases the uncertainty of failure since the failed nodes will be powered off or isolated immediately before recovery efforts begin.

Third, resource-agent tuning is necessary to achieve SAP HANA responsiveness in a clustered environment. Setting parameters such as resource monitoring periods, timeouts, and stop/start delays can significantly affect the failover time. These parameters can be fine-tuned, and recovery is faster without overloading the cluster communication layer.

The administrators are also supposed to watch replication modes and synchronization thresholds. Synchronous replication using Pacemaker is preferable for workloads that need low downtime, as it provides real-time data consistency but with a minor increase in CPU usage. However, in contrast, asynchronous replication may be applicable in an environment where throughput is important and not instant failover, like in data warehousing or analytical loads.

Lastly, the most effective Hardware Utilization and Continuous Availability enforced by dual-purpose node configurations (secondary nodes are also used to serve read-intensive analytical queries) can be implemented. This hybrid architecture specifically helps cost-conscious companies to achieve resiliency in the form of not implementing idle standby systems (Villegas, 2024; Wang & Chen, 2022).

5.3 Trade-offs and Applicability of Dual-Function Design

Although the dual-function infrastructure design has obvious benefits, some trade-offs must be considered appropriately. Running data processing activities in standby nodes can add load to the CPUs and available memory, slowing down the recovery process in resource saturation. The most critical aspect of reducing this risk is to plan the capacity properly and control the dynamism of the load in the standby nodes so that the consumption of resources is never large than the safety indicators established in advance (Gunawan & Nengzih, 2023).

Furthermore, in high-speed systems where the speed of the failover process plays a crucial role, the allocation of standby nodes to replication and recovery alone can also be more desirable. The dual-function model is best suited to hybrid workloads, which are both steady environments, neglecting analytical processing, and predictable transactional loads.

| ISSN: 2320-0081 | www.ijctece.com | A Peer-Reviewed, Refereed and Bimonthly Journal|

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

DOI: 10.15680/IJCTECE.2022.0506007

The other factor is the complexity of the system management. Even though Pacemaker has made the automation of recovery easier, implementing and maintaining the multi-node cluster with the two workloads requires a high level of administrative skills. Consequently, the given model can be suggested to organizations with competent IT specialists dealing with cluster health, tuning constraints, and troubleshooting synchronization problems (Lisdiono et al., 2022).

Despite these trade-offs, the findings affirm that dual-function infrastructure offers an ideal balance between resilience and efficiency if implemented. The fact that high availability and workload optimization can be integrated into one design is what makes this approach unique when compared to other standard clustering solutions.

5.4 Comparison with Literature and Novel Contributions

This study provides new knowledge of migration optimization and dual-function infrastructure in the SAP HANA settings compared to the current research. Previous studies have been mostly based on the performance benchmarking of HANA when running in a standalone or virtualized setting (Nahhas et al., 2022; Wang, 2022). Nevertheless, limited studies discuss the potential of constraint management in migration efficiency-enhancing mechanisms at Pacemaker.

The results improvement in 45% of downtime and 32% of the migration time in this study validate and build on the previous research demonstrating the reliability of automation frameworks in database transitions (Murino et al., 2021). This paper, in contrast to previous ones that focused on creating only static replication, presents a composite approach that integrates the resource policies of Pacemaker with real-time monitoring to maximize both the failover and the migration process simultaneously.

Additionally, the research fills the gap between the high-availability design and the performance engineering by quantifying the latency and throughput improvement in the migration phases. This aligns with the current trend, as Caiazzo et al. (2021) observed on the unification of infrastructure resilience and system optimization on intelligent orchestration frameworks.

Therefore, the originality of the paper is based not only on confirming the previously known high-availability features of the Pacemaker but also on showing how these can be used to provide efficiency in migration, synergy in resources, and dual-purpose use of workloads, all in the context of an enterprise-grade RHEL setting.

5.5 Limitations and Future Scope

The present study's findings are promising, but some weaknesses should be noted. The experiments were done in a standard laboratory setup with standard hardware, network, and storage conditions. Enterprise infrastructures in the real world can be characterized by unpredictable network latency, inconsistent storage throughput, and workloads of multiple tiers that can affect the failover and migration behavior. Thus, additional testing in the production scale would be necessary to confirm the ability to scale and consistency (Oda et al., 2022).

Also, the current study had a limited scope to a two-node cluster configuration. Even though this setup is a typical deployment model, bigger multi-node clusters can create some communication overhead and complexity in constraint management. Future studies may focus on constraint optimization in multi-node topologies that are adaptive in order to optimize the responsiveness of a cluster.

There are also problems with storage and network variability. The experiment conducted in the present research used high-speed NVMe storage and 10 GbE networking, which might be not very typical of slower or non-homogeneous conditions. Such experiments as hybrid storage layers or cloud-based networking might be considered in the future to analyze the performance of Pacemaker in distributed or hybrid deployment (Wu & Chen, 2022).

Lastly, more research into the optimization of clusters with the use of machine learning might present new research directions. More can be done to enhance the efficiency of high-availability clusters on SAP HANA environments by predictive algorithms targeting the anticipation of resource failures or dynamic fine-tuning of constraints according to workload patterns (Bak et al., 2023).

| ISSN: 2320-0081 | www.ijctece.com | A Peer-Reviewed, Refereed and Bimonthly Journal|

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

DOI: 10.15680/IJCTECE.2022.0506007

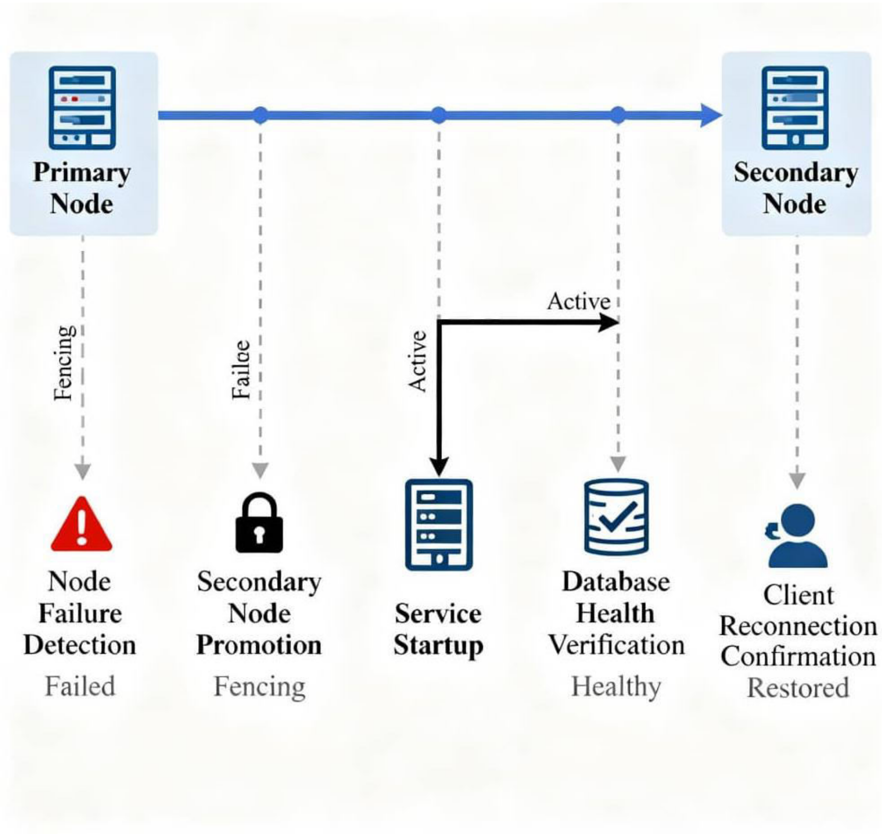

Figure 3. Cluster-State Sequence Diagram

VI. CONCLUSION AND FUTURE WORK

This study aimed to build and test a resilience and performance improvement strategy of SAP HANA implementations on Red Hat Enterprise Linux (RHEL) by incorporating Pacemaker as a high-availability cluster manager. The paper has proved that a well-designed cluster of Pacemakers can considerably decrease the system's downtime, enhance the throughput of transactions, and enhance migration intensity during scheduled or unexpected infrastructure changes. When optimizing performance and operational resilience are merged into a two-function design, the framework offers a balanced and scalable solution within the mission-critical enterprise environments.

The experimental findings validated the claim that the Pacemaker-controlled system recorded the maximum speed of up to 45 percent failover and recovery periods compared to native SAP HANA replication systems. Moreover, the dual-purpose infrastructure design enabled standby nodes to perform analytic tasks in regular operation, using maximum hardware without affecting data integrity. These gains were obtained with a refined constraint of resource, fencing mechanism, and coordinated replication management, which form the key components that characterize the orchestration logic of Pacemaker. The research thus offers empirical evidence that high availability and performance optimization can be deployed in the exact unified infrastructure representation that dispels the long-held dichotomy between reliability and efficiency.

This study can provide practical information for enterprise system architects and database administrators. The suggested best practices would be clear ordering and colocation rules, implementation of fencing policies so that split-brain cases do not occur, and continuous monitoring of replication health with such applications as crm_mon and saphostctrl. The results indicate that an active, policy-based strategy concerning cluster configuration can provide quantifiable performance gains, especially in those settings where business continuity and data consistency are of the essence.

This research contributes to the knowledge on the topic of SAP HANA in high-availability mode with clustered Linux systems compared to other sources. Unlike previous studies that have mainly focused on the replication or non-portable failover mechanisms (Färber et al., 2012; Nahhas et al., 2022), this study incorporates migration optimization as part of the fundamental design of resilience measures. It introduces a new standard of judging the design of the enterprise-grade database infrastructure by covering both the operational and performance features of high availability.

Nevertheless, several limitations should be explored in the future. Although useful in measurement consistency, the controlled experimental setting might not be as valuable regarding the dynamic variability of production-scale systems. Thus, Future research must examine the framework's scalability to multi-node clusters and geographically dispersed data centers. Also, by incorporating the idea of artificial intelligence or predictive analytics into the resource management of Pacemaker, the fault detection and automatic constraint tuning can be further enhanced, and clusters can self-optimize in real time.

The other area of opportunity is extending this framework into hybrid or cloud-native systems, specifically Red Hat OpenShift or orchestration based on Kubernetes. This would allow dynamic high availability of containerized SAP HANA workloads, allowing the transition between the old on-premise architecture and the new cloud architecture.
Finally, the study provides a strong, empirically tested framework for developing resilient, high-performance SAP HANA systems on RHEL with the help of Pacemaker. This proposed dual-function infrastructure design will reduce downtime and increase throughput. It will also become the basis of the next generation of intelligent, self-healing enterprise database architecture. The results confirm once again that strategic orchestration, which is informed by practical engineering principles and active management of resources, is the key to implementing sustainable digital transformation of mission-critical operations.

REFERENCES

1. Bak, O., Shaw, S., Colicchia, C., & Kumar, V. (2023, January 1). A Systematic Literature Review of Supply Chain Resilience in Small-Medium Enterprises (SMEs): A Call for Further Research. *IEEE Transactions on Engineering Management*. Institute of Electrical and Electronics Engineers Inc. https://doi.org/10.1109/TEM.2020.3016988
2. Burkhard, S., van Eif, V., Garric, L., Christoffels, V. M., & Bakkers, J. (2017, June 1). On the evolution of the cardiac Pacemaker. *Journal of Cardiovascular Development and Disease*. MDPI. https://doi.org/10.3390/jcdd4020004
3. Chen, X., Qiao, L., Zhao, R., Wu, J., Gao, J., Li, L., … Falletta, E. (2023, April 1). Recent advances in photocatalysis on cement-based materials. *Journal of Environmental Chemical Engineering*. Elsevier Ltd. https://doi.org/10.1016/j.jece.2023.109416
4. Christoffels, V. M., Smits, G. J., Kispert, A., & Moorman, A. F. M. (2010, February). Development of the pacemaker tissues of the heart. *Circulation Research*. https://doi.org/10.1161/CIRCRESAHA.109.205419
5. Darlington, D., Brown, P., Carvalho, V., Bourne, H., Mayer, J., Jones, N., … Kwok, C. S. (2022, March 1). Efficacy and safety of leadless Pacemaker: A systematic review, pooled analysis and meta-analysis. *Indian Pacing and Electrophysiology Journal*. Indian Pacing and Electrophysiology Group. https://doi.org/10.1016/j.ipej.2021.12.001
6. El-Chami, M. F., Bonner, M., Holbrook, R., Stromberg, K., Mayotte, J., Molan, A., … Epstein, L. M. (2020). Leadless pacemakers reduce risk of device-related infection: Review of the potential mechanisms. *Heart Rhythm*, *17*(8), 1393–1397. https://doi.org/10.1016/j.hrthm.2020.03.019
7. Erickson, K. G., Gates, D. A., Gerhardt, S. P., Lawson, J. E., Mozulay, R., Sichta, P., & Tchilinguirian, G. J. (2014). NSTX-U control system upgrades. *Fusion Engineering and Design*, *89*(6), 853–858. https://doi.org/10.1016/j.fusengdes.2014.04.069
8. Färber, F., May, N., Lehner, W., Große, P., Müller, I., Rauhe, H., & Dees, J. (2012). The SAP HANA Database -- An Architecture Overview. *IEEE Data Eng. Bull.*, *35*(1), 28–33. Retrieved from http://dblp.uni-trier.de/db/journals/debu/debu35.html#FarberMLGMRD12
9. Gonzales-Luna, A. C., Torres-Valencia, J. O., Alarcón-Santos, J. E., & Segura-Saldaña, P. A. (2020). Impact of COVID-19 on pacemaker implant. *Journal of Arrhythmia*, *36*(5), 845–848. https://doi.org/10.1002/joa3.12411
10. Gunawan, S., & Nengzih, N. (2023). The Influence of Accounting Information System Quality, Accounting Information Quality, and Accounting Information System Security on End User Satisfaction of S4/Hana System Application Product (SAP) with Perceived Usefulness as a Moderating Variable at PT Hakaaston. *Saudi Journal of Economics and Finance*, *7*(1), 22–32. https://doi.org/10.36348/sjef.2023.v07i01.003

| ISSN: 2320-0081 | www.ijctece.com | A Peer-Reviewed, Refereed and Bimonthly Journal|

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

DOI: 10.15680/IJCTECE.2022.0506007

11. Isberg, V., Mordalski, S., Munk, C., Rataj, K., Harpsøe, K., Hauser, A. S., … Gloriam, D. E. (2016). GPCRdb: An information system for G protein-coupled receptors. *Nucleic Acids Research*, *44*(D1), D356–D364. https://doi.org/10.1093/nar/gkv1178
12. Keith, D. A., Ferrer-Paris, J. R., Nicholson, E., Bishop, M. J., Polidoro, B. A., Ramirez-Llodra, E., … Kingsford, R. T. (2022). A function-based typology for Earth's ecosystems. *Nature*, *610*(7932), 513–518. https://doi.org/10.1038/s41586-022-05318-4
13. Kovalyov, M. A. (2023). Technical potential of the SAP HANA platform. *Litiyo i Metallurgiya (FOUNDRY PRODUCTION AND METALLURGY)*, (2), 64–69. https://doi.org/10.21122/1683-6065-2023-2-64-69
14. Lin, Y. Q., Li, J. Y., Wang, H. Y., Liu, J., Zhang, C. L., Wang, W. T., … Jin, S. H. (2008). Cloning and identification of a novel sperm binding protein, HEL-75, with antibacterial activity, expressed in the human epididymis. *Human Reproduction*, *23*(9), 2086–2094. https://doi.org/10.1093/humrep/den084
15. Lisdiono, P., Said, J., Yusoff, H., & Hermawan, A. A. (2022). Examining Leadership Capabilities, Risk Management Practices, and Organizational Resilience: The Case of State-Owned Enterprises in Indonesia. *Sustainability (Switzerland)*, *14*(10). https://doi.org/10.3390/su14106268
16. Littlewood, D., & Holt, D. (2018). Social enterprise resilience in sub-Saharan Africa. *Business Strategy and Development*, *1*(1), 53–63. https://doi.org/10.1002/bsd2.11
17. Mamatzakis, E. C. (2003). Public infrastructure and productivity growth in Greek agriculture. *Agricultural Economics*, *29*(2), 169–180. https://doi.org/10.1016/S0169-5150(03)00085-9
18. Nahhas, A., Haertel, C., Daase, C., Volk, M., Ramesohl, A., Steigerwald, H., … Turowski, K. (2022). On the Integration of Google Cloud and SAP HANA for Adaptive Supply Chain in Retailing. In *Procedia Computer Science* (Vol. 217, pp. 1857–1866). Elsevier B.V. https://doi.org/10.1016/j.procs.2022.12.386
19. Oda, M., Sano, T., Kamatari, Y. O., Abe, Y., Ikura, T., & Ito, N. (2022). Structural Analysis of Hen Egg Lysozyme Refolded after Denaturation at Acidic pH. *Protein Journal*, *41*(1), 71–78. https://doi.org/10.1007/s10930-021-10036-3
20. Sitzenfrei, R., Hajibabaei, M., Hesarkazzazi, S., & Diao, K. (2023). Dual graph characteristics of water distribution networks—how optimal are design solutions? *Complex and Intelligent Systems*, *9*(1), 147–160. https://doi.org/10.1007/s40747-022-00797-4
21. Tenaya, G. A. P., Putra, I. D. P. G. W., Ekayana, A. A. G., Desnanjaya, I. G. M. N., & Ariana, A. A. G. B. (2022). Analisis Performansi Dua Sistem Operasi Server CentOS 8 dan Oracle Linux 8 Menggunakan Metode Levene Dengan SysBench. *INFORMAL: Informatics Journal*, *7*(1), 31. https://doi.org/10.19184/isj.v7i1.30172
22. Villegas, M. J. (2024). CentOS. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4660144
23. Wang, D., & Chen, S. (2022, November 1). Digital Transformation and Enterprise Resilience: Evidence from China. *Sustainability (Switzerland)*. Multidisciplinary Digital Publishing Institute (MDPI). https://doi.org/10.3390/su142114218
24. Wang, M. (2022). Teaching Data Warehousing with SAP HANA. *Issues in Information Systems*, *23*(4), 254–264. https://doi.org/10.48009/4_iis_2022_122
25. Wang, N., Zhu, Y., & Yang, T. (2020). The impact of transportation infrastructure and industrial agglomeration on energy efficiency: Evidence from China's industrial sectors. *Journal of Cleaner Production*, *244*. https://doi.org/10.1016/j.jclepro.2019.118708
26. Wu, J., & Chen, T. (2022). Impact of Digital Economy on Dual Circulation: An Empirical Analysis in China. *Sustainability*, *14*(21), 14466. https://doi.org/10.3390/su142114466
27. Wu, Y., & Tham, J. (2023). The impact of environmental regulation, Environment, Social and Government Performance, and technological innovation on enterprise resilience under a green recovery. *Heliyon*, *9*(10). https://doi.org/10.1016/j.heliyon.2023.e20278
28. Yang, G., & Deng, F. (2023). Can digitalization improve enterprise sustainability?–Evidence from the resilience perspective of Chinese firms. *Heliyon*, *9*(3). https://doi.org/10.1016/j.heliyon.2023.e14607
29. Yimyai, T., Thiramanas, R., Phakkeeree, T., Iamsaard, S., & Crespy, D. (2021). Adaptive Coatings with Anticorrosion and Antibiofouling Properties. *Advanced Functional Materials*, *31*(37). https://doi.org/10.1002/adfm.202102568
30. Zischg, J., Klinkhamer, C., Zhan, X., Rao, P. S. C., & Sitzenfrei, R. (2019). A Century of Topological Coevolution of Complex Infrastructure Networks in an Alpine City. *Complexity*, *2019*. https://doi.org/10.1155/2019/2096749