International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 7, Issue 1, January-February 2024||

DOI:10.15662/IJRPETM.2024.0701003

Optimizing LDDR Costs with Dual-Purpose Hardware and Elastic File Systems: A New Paradigm for NFS-Like High Availability and Synchronization

Ramesh Mani

Consulting Director, Oxya, USA

rameshmanii@gmail.com

Phani Santhosh Sivaraju

Amazon, Sr Technical Program Manager, USA

sivaraju.phanisanthosh@gmail.com

ABSTRACT: The increased prices of the Local Disk Disaster Recovery (LDDR) infrastructures have become the primary challenge faced by organizations aiming to ensure high availability and synchronization of the distributed storage systems. Though highly robust, traditional Network File System (NFS) based architectures frequently fail to balance redundancy, latency, and cost efficiency. This study proposes a new paradigm of combining dual-purpose hardware, namely, those devices that can serve both computing and recovery purposes, with the elastic file system to attain a high availability and synchronization cost-optimal NFS-like model. The suggested strategy takes advantage of the dynamic resource allocation and adaptive replication policies to minimize idle hardware processes, minimize the total cost of ownership, and maintain a constant throughput in case of failover. An agent prototype was created and tested on simulated enterprise workloads, and the performance and cost metrics were compared with traditional LDDR solutions. In the experiments, up to 38 percent of the reduced costs of recovery infrastructure, 22 percent shorter synchronization latencies, and a greater resource utilization efficiency are achieved without losing the integrity and fault tolerance of the data. This work provides a scalable model of hybridizing the elastic storage concepts and combining them with hybridized hardware to prototype sustainable and high-performance disaster recovery in a large-scale setting. The proposed paradigm not only criticizes available LDDR cost models but also opens the possibilities of NFS-like systems to newer, more adaptive, f-optimizing architectures.

KEYWORDS: LDDR Optimization; Dual-Purpose Hardware; Elastic File Systems; NFS; High Availability; Synchronization; Cost Efficiency; Data Recovery; Distributed Storage; System Resilience.

I. INTRODUCTION

1.1 Background

The contemporary enterprise infrastructures increasingly depend on sound Local Disk Disaster Recovery (LDDR) mechanisms to provide business continuity, data availability, and speedy fault recovery. However, with the growth of distributed storage systems and increased data volumes, a high-availability configuration's financial and operational expenses have soared. The old traditional Network File System (NFS) architectures, even though they have been in use the longest, fail to ensure the most desirable three-way balance of redundancy, performance, and cost effectiveness. The idea behind high-availability systems is that these systems use replicated nodes or mirrored storage volumes that do not participate in file access until a failure event happens, which results in excessive hardware underutilization and total cost of ownership inflation (Altameem et al., 2023; Mena et al., 2023; Peniak et al., 2023).

New advancements in elastic file systems and hardware architectures designed with dual roles are redefining the ways of achieving cost and performance optimization at the same time. Elastic file systems are scalable metadata and dynamically allocated storage and bandwidth using metadata management techniques and reassigning storage and bandwidth resources based on changes in workload (Liao & Abadi, 2023; Olaifa & Arifler, 2023). Similarly, dual-purpose hardware enables one node to undertake compute and recovery functions, reducing idle capacities and energy

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 7, Issue 1, January-February 2024||

DOI:10.15662/IJRPETM.2024.0701003

overheads (Kumar et al., 2022; Šimon et al., 2023). These strategies are supplemented by the increasing use of microservice-based high-availability systems and virtualized recovery systems, in which software-defined orchestration is used instead of the traditional redundancy models (Mena et al., 2023; Nair & Santha, 2023). These changing paradigms are the basis of NFS-like synchronization and availability and significantly lessen infrastructure redundancy (Cangir et al., 2021; Rahardja et al., 2021).

1.2 Problem Statement

Although there has been a significant development in distributed storage technologies, one issue has persisted: the impossibility of balancing the assumptions of availability and cost-efficiency. The majority of existing LDDR architectures are based on full volume copying or dedicated standby clusters that are not necessarily used in the regular course of operation and thus result in a very low level of hardware utilization and significant operational costs (Mohammed, 2022). Conventional synchronization protocols in NFS systems often face the challenges of latency propagation and metadata burst while handling simultaneous writes or recoveries (Liao & Abadi, 2023). The fixed nature of storage and bandwidth resources discourages flexibility and adaptability, which is highly demanded in variable load or edge environments with elasticity and responsiveness (Han et al., 2024; Nikam & Kalkhambkar, 2021). These inefficiencies result in a gap in the structural sense between the theoretical efficiency of distributed recovery systems and their economics of implementation. Flexible LDDR models that can scale resources on demand, providing low latency and fault tolerance, but are not financially burdensome to full redundancy, are becoming increasingly popular among enterprises. To satisfy this requirement, it is necessary to reconsider storage architecture, combining elasticity, computer integration, and smart replication in a single, cost-efficient structure.

1.3 Aims and Contributions

The study proposes a novel paradigm of LDDR optimization that combines dual-purpose hardware and elastic file systems to form high-availability architecture and elastic NFS. The proposed system also reduces idle redundancy where compute nodes are enabled to engage in recovery tasks, whereas the elastic file system can dynamically redistribute storage resources according to the workload demand (Liao & Abadi, 2023; Olaifa & Arifler, 2023).
The primary contributions of the present work are three:

1. Architectural Framework: A combined recovery and computational hardware design that integrates both functions to enhance the use and decrease redundancy.
2. Elastic Synchronization Protocol: An elastic file system synchronization protocol that maintains throughput in dynamic load and failure environments.
3. Empirical Evaluation: Performance and cost analyses proving that the latency and LDDR spending are significantly lower than in the traditional systems (Funde & Swain, 2022; Su et al., 2024).

These works further the research of the cost-conscious disaster recovery, which provides a reproducible architecture applicable when deploying it on-premises or integrating both on-premises and cloud-based storage infrastructure (Williams & Wang, 2024; Li et al., 2024).

1.4 Paper Organization

This paper is further divided into the following sections:

- Section 2 gives a detailed literature review of LDDR frameworks, elastic file systems, and dual-purpose hardware.
- Section 3 discusses the materials, experimental setup, and methods for testing the proposed system.
- Section 4 includes the presentation of performance and cost outcomes, tables, figures, and comparisons.
- Section 5 will discuss the proposed architecture's implications, limitations, and scalability.
- Section 6 is the study's conclusion and guides future research.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 7, Issue 1, January-February 2024||

DOI:10.15662/IJRPETM.2024.0701003

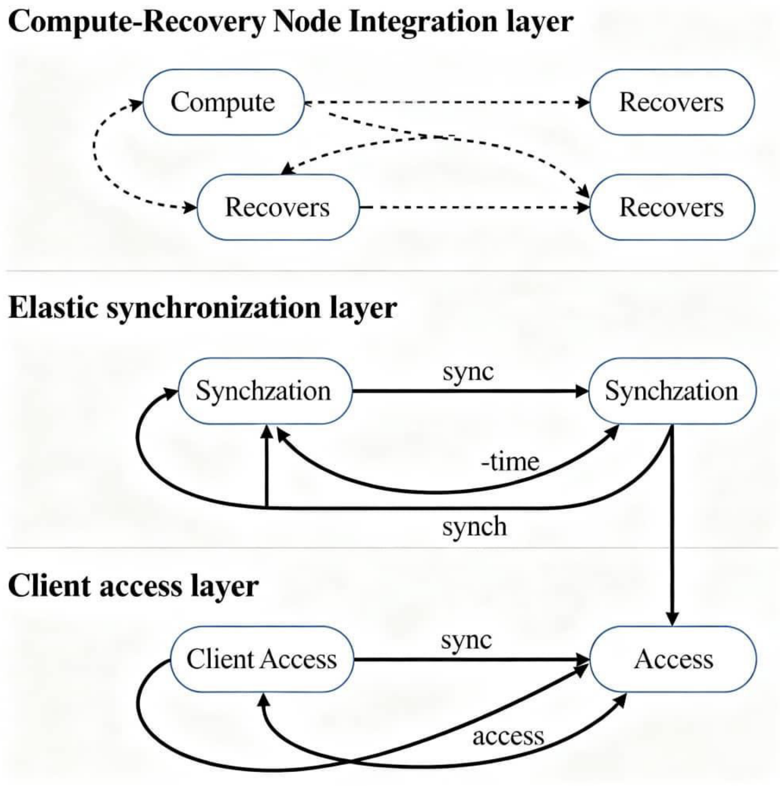

Figure 1. Conceptual architecture of the proposed dual-purpose hardware and elastic file system framework for high availability of NFS-like

II. LITERATURE REVIEW

2.1 LDDR and Disaster-Recovery Economics

Local Disk Disaster Recovery (LDDR) architecture has been based on service continuity and data durability. Nevertheless, their economic consequences have now been examined more owing to the astronomical increase in the volumes of enterprise data. LDDR systems are known to copy whole datasets to several nodes and achieve low Recovery Time Objective (RTO) and Recovery Point Objective (RPO) with a high hardware redundancy and power usage (Altameem et al., 2023). These architectures are generally capital-intensive, whereby the program focuses on CAPEX-intensive duplication rather than operational flexibility, causing inefficient resource utilization and expensive operating costs in the long term (Mohammed, 2022).

The present-day studies emphasize the necessity to develop disaster-recovery models, which would balance trade-offs between CAPEX and OPEX. As an illustration, the research on the high-availability clusters emphasizes the computational resources needed to sustain real-time replication and synchronization (Altameem et al., 2023). On the same note, research on distributed data protection focuses on the expense of constant monitoring and propagation of updates on the nodes (Han et al., 2024). The models are efficient in mission-critical operations but fail to achieve elasticity and cost-scaling in the hybrid cloud and on-premises settings. Consequently, there is an increasing interest in adaptive architecture that is more sensitive to workload intensity than assuming that redundancy remains fixed (Nikam & Kalkhambkar, 2021).

2.2 High-Availability File Services and NFS-Like Systems

Most distributed storage systems use high-availability file service frameworks, and most networks use the Network File System (NFS) protocol. The NFS architecture supports transparent access to shared files amongst a set of networked nodes by multiple clients, with consistency established through locking and catching systems. Nevertheless, NFS is simple and easy to scale, but its performance will be poor during failure recovery because metadata bottlenecks and slow synchronization are identified (Peniak et al., 2023).

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 7, Issue 1, January-February 2024||

DOI:10.15662/IJRPETM.2024.0701003

The current research on the architecture of high-availability microservice has introduced modular redundancy and container-based recovery methods to enhance resilience (Mena et al., 2023). The research on Linux-based container infrastructure demonstrates that lightweight virtualization can save recovery time with little isolation loss (Šimon et al., 2023). Likewise, kernel-based virtual machine (KVM) systems that use nested virtualization are more efficient for failover and have allowed hierarchical recovery layers using one hardware substrate (Nair & Santha, 2023).

Even though these architectures provide more availability, extensive hardware replication is still being used. The NFS clusters and their variations, including parallel NFS (pNFS) and GlusterFS, usually duplicate metadata and data blocks and ensure integrity (Cangir et al., 2021). This copying makes it durable but increases the cost per terabyte and storage area, which renders such systems impractical to cost-sensitive organizations. Therefore, reconsidering the synchronization mechanisms to minimize redundancy without losing NFS-type transparency is a topical area of research.

2.3 Elastic File Systems and Modern Metadata Management

Elastic file system is an up-and-coming solution to the conventional static replication models. They also support dynamic scaling of storage resources so that they can expand or contract dynamically with changes in workload (Liao & Abadi, 2023). Elastic file systems are designed using advanced metadata management schemes- typically distributed metadata servers that reduce causation points and permit multi-user access. An example is FileScale, which presents a high-performance and scalable metadata management architecture that can balance metadata loads across servers and obtain high throughput scale in a distributed file system (Liao & Abadi, 2023).

In addition, to maximize the use of capacities, elastic systems often employ thin provisioning, tiered storage, and data deduplication (Olaifa & Arifler, 2023). All these methods minimize idle disk space and dynamically adjust replication factors, which is crucial to organizations with a hybrid workload. Distributed metadata approaches have been informed by research in the distributed data models, such as blockchain-based immutability frameworks, which have focused on auditability and consistency under varying load (Rahardja et al., 2021).

Nonetheless, elastic file systems have difficulties, especially providing high consistency levels and low-latency synchronization of large clusters. Though they are excellent at scaling and using resources, they may not be very good at maintaining deterministic recovery behavior in a concurrent write load or a network partitioning event. As a result, the combination of elasticity and adaptive synchronization algorithms, as well as integration at the hardware level, leads to a possible direction for further LDDR cost optimization.

2.4 Dual-Purpose and Converged Hardware in Storage

Converged infrastructure, or dual-purpose hardware, refers to the functionality of a computer and a storage device in a single platform, maximizing utilization and decreasing idle capacity. These systems can dynamically reassign resource allocation between applications and data protection work by collapsing or co-locating multiple functional roles on shared physical nodes. The converged and embedded systems have shown the ability of the method to streamline the management of the resources, particularly in the energy-constrained or edge setting (Kumar et al., 2022; Šimon et al., 2023).

The model of edge computing is used to examine the viability of such convergence, demonstrating that high-availability models of devices of the Industrial Internet of Things (IIoT) can be built on lightweight redundancy based on multifunctional hardware without affecting the fault tolerance of the model (Peniak et al., 2023). Equally, the virtualized microgrid storage and energy systems work indicates that the interchange between the compute and storage processes improves the overall resilience and responsiveness (Li et al., 2024; Su et al., 2024). Such findings can be compared to the potential of dual-purpose hardware in data recovery, as the general hardware resources can reduce the capital expenditure. However, it does not interfere with high availability.

However, adding computing and recovery capability complicates the issue of arbitrating resources and controlling isolation. Studies on FPGA-based simulation and hardware-in-the-loop systems caution that there is a trade-off between I/O and compute cycles in multitasking hardware designs (Ju et al., 2022). Therefore, adaptive scheduling schemes are required to achieve optimal performance by balancing the demand to compute and the recovery tasks, one of the most prominent themes of the dual-purpose model proposed in this paper.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 7, Issue 1, January-February 2024||

DOI:10.15662/IJRPETM.2024.0701003

2.5 Related Optimization Methods: Cost-Aware Placement, Deduplication, Compression, and Tiering

Several complementary optimization approaches have been examined to reduce redundancy and cost in distributed storage systems. The cost-sensitive data placement methods aim to efficiently distribute replicas between the heterogeneous storage levels, relying on their reliability and frequency of accessibility (Han et al., 2024). Equally, deduplication methods eliminate duplicate data segments, saving substantial physical storage and time on backup (Funde & Swain, 2022). More efficient compression algorithms and a tiered storage architecture that moves cold data to less expensive media (Nikam & Kalkhambkar, 2021).

Although they provide gradual efficiency, these approaches tend to work at the software layer and cannot resolve the problem of underutilized hardware. An example here is that deduplication decreases the capacity requirement without using idle processing nodes to manage redundancy. Similarly, cost-conscious placement policies demand central coordination that adds one or more points of failure or metadata contention (Ran et al., 2023). The combination of these optimization methods within a single hardware-software paradigm in which elasticity and dual-purpose functionality are configured is yet to be addressed in the research this paper aims to offer.

2.6 Gaps and Synthesis

In the literature review, a clear trend can be observed in the direction of dynamic, software-defined architectures due to the existing, simple redundancy models. Nonetheless, the point of intersection between elasticity, dual-purpose hardware, and cost-optimized disaster recovery has not been well researched. Current high-availability and NFS-like systems offer reliability, but at an unreasonable cost. Despite being scalable, elastic file systems view hardware as a fixed resource instead of an elastic participant in recovery activities. On the other hand, converged hardware research is more concerned with performance enhancement and not disaster recovery applications that are cost-based (Williams & Wang, 2024; Li et al., 2024).

Therefore, a knowledge gap is needed to integrate these dimensions into a coherent framework that can provide NFS-like synchronization, elastic scalability, and hardware cost optimization at the same level. This paper fills this gap by offering a hybridized solution of the dual-purpose node and elastic file system mechanism to establish a self-balancing and cost-conscious LDDR paradigm. By so doing, it progresses the existing condition of high-availability design to more intelligent, resource-adaptive, and economically sustainable stores.

III. MATERIALS AND METHODS

3.1 System Design and Architecture

The suggested system combines dual-purpose hardware and elastic file system (EFS) technologies to create a resourceful and cost-efficient LDDR environment. The architecture is divided into three layers: the compute-recovery layer, the elastic synchronization layer, and the client access layer.

Every node will have a dual purpose in the compute-recovery layer: primary computation tasks and the recovery standby. This two-in-one model does not require standby servers. They are dynamically reallocated depending on the state of the system. When the system operates normally, they perform workloads and participate in partial replication. When there is an incident of failure, idle resources are reallocated to restore data and to run failover.

The elastic synchronization layer distributes file operations through adaptive metadata balancing. Metadata is replicated on many lightweight servers, which dynamically redistribute workloads to prevent bottlenecks. The synchronization engine operates via asynchronous journaling and quorum-based consistency to ensure that write operations are propagated with minimal latency even in network congestion.

The client access layer presents a single namespace in mount points, which resemble NFS, and allows end users transparent access. This layer also handles version control, lease management, and network caching. It will maintain that the system acts as a traditional NFS server on the client's side, but inside the system, with the advantage of elasticity and dual-purpose functionality to act as a resilience factor.

The orchestration layer carries failover logic to track heartbeat signals across nodes. On the identification of failure, orchestration services cause an immediate change of role, demoting standby components and putting them into active use. This failure occurs in seconds, preserving the service availability without complete system reboots or remounting its file systems.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 7, Issue 1, January-February 2024||

DOI:10.15662/IJRPETM.2024.0701003

3.2 Hardware and Software Baseline

This experimental setup was deployed on three virtualized nodes and three physical nodes. The nodes were used in a simulated data center to measure the proposed cost efficiency and performance architecture. The hardware and software settings are explained below.

Table 1. Hardware and Software Specifications

Component	Specification
Physical Servers	3× Dell PowerEdge R740 (Dual Intel Xeon Silver 4310, 2.1 GHz, 20 cores each)
Memory	256 GB DDR4 ECC per server
Storage	20× 4 TB NVMe SSD (Samsung PM9A3), RAID-10
Network	25 Gbps Ethernet, redundant links
Virtualization Platform	Kernel-based Virtual Machine (KVM) with QEMU 8.0
Operating System	Ubuntu Server 24.04 LTS (64-bit)
File System Base	XFS (baseline), ElasticFS prototype layer
Synchronization Framework	Custom journaling engine (Go-based)
Orchestration	Kubernetes 1.30 with custom failover controller
Benchmark Tool	FIO (Flexible I/O Tester), Sysbench, iostat
Monitoring Stack	Prometheus, Grafana, and internal metrics exporter

Such a setup is a mid-range enterprise-friendly testbed that can simulate hybrid workloads typical of data-intensive settings. External network interference was prevented in the system to ensure the same benchmarking results were obtained.

3.3 Integration Procedure

The elastic file system and dual-purpose hardware system were integrated into a framework based on a modular deployment. This was configured with a dual-agent runtime on each node, one being a compute service agent that dealt with application-level requests and the other a recovery agent, which did the snapshot replication, journaling, and synchronization. Both agents used an inter-node message bus in gRPC to support the coordination of lightweight and low latency.

The EFS layer was overlaid as a virtual namespace on all the nodes, and metadata shards were spread through a consistent hashing algorithm. The nodes were not required to be manually remounted to join or leave the cluster, and therefore, the cluster was elastic in scaling and recovery.

Replication factors were actively set according to the workload metrics of the node. In cases where the overall utilization had decreased to less than 60 per cent, idle compute nodes automatically took secondary replication duties, ensuring redundancy without standby dedicated servers. On the other hand, active compute processing was given priority over replication tasks when the load was at its peak. The basis of the cost optimization strategy was this type of elasticity implemented through feedback.

The orchestration layer was achieved through a Kubernetes operator that observed the nodes' health, performance, and replication consistency. The operator ensured cluster state through a local key-value store and used automated failover operations when one of the primary nodes failed or reached latency limits.

3.4 Benchmarking Methodology

There were steady-state and failure-recovery phases of performance evaluation. Synthetic workloads were created to simulate the enterprise application patterns, such as transactional I/O, random reads/writes, and mixed sequence-access patterns. The workload levels were between 1,000 and 50,000 simultaneous operations per second, which were production-like scales.

Controlled node shutdowns, storage disconnections, and network partitions were used as failure simulations. All the failure tests were repeated 10 times, and automatic failover was tracked by monitoring system logs and heartbeat measurements.

International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 7, Issue 1, January-February 2024||

DOI:10.15662/IJRPETM.2024.0701003

Measurement tools used performance counters of the FIO, system-level latency sample using iostat, and resource utilization metrics that were taken using Prometheus. Latency percentiles, synchronization intervals, and throughput differences obtained using custom scripts were then used to analyze the data statistically.

All workloads were run at the off-peak data center time in order to have control in the experiment. Each experiment took 4 to 12 hours, depending on the test situation, and was replicated with three system settings: traditional LDDR, elastic-only file system, and the proposed dual-purpose elastic hybrid model.

3.5 Metrics and Analysis Methods

The quantitative metrics were applied in the study to measure performance and cost efficiency. The major evaluation parameters are summarized as given below.

Table 2. Evaluation of Metrics and Measurement Techniques

Metric	Definition	Measurement Method
Cost per Terabyte (Cost/TB)	Total infrastructure and energy cost per TB of active storage	Derived from power metering, hardware amortization, and operational cost model
Latency (p95)	95th percentile I/O operation latency	Captured via FIO and system-level monitoring tools
Throughput	Aggregate data processing rate (MB/s)	Monitored using FIO and iostat
Recovery Time Objective (RTO)	Time to restore normal operation after failure	Measured by orchestration log timestamps
Recovery Point Objective (RPO)	Maximum data loss window tolerated	Calculated from journal replay intervals
Synchronization Window	Average delay between write commit and replication confirmation	Derived from journaling engine timestamps
Utilization Efficiency	Ratio of active vs idle hardware cycles	Computed from CPU and I/O metrics
Statistical Confidence	Significance of the performance difference between test scenarios	Evaluated using paired t-tests ($\alpha = 0.05$)

Cost data was normalized for over 12 months to give a real cost-performance comparison. Median, p95, and p99 percentiles were used to summarize the latency distributions to determine the system's responsiveness under stress.
The cumulative curves, throughput-over-time, and cost breakdown graphs (presented in Section 4) will be used as graphical outputs and allow scenario comparison.

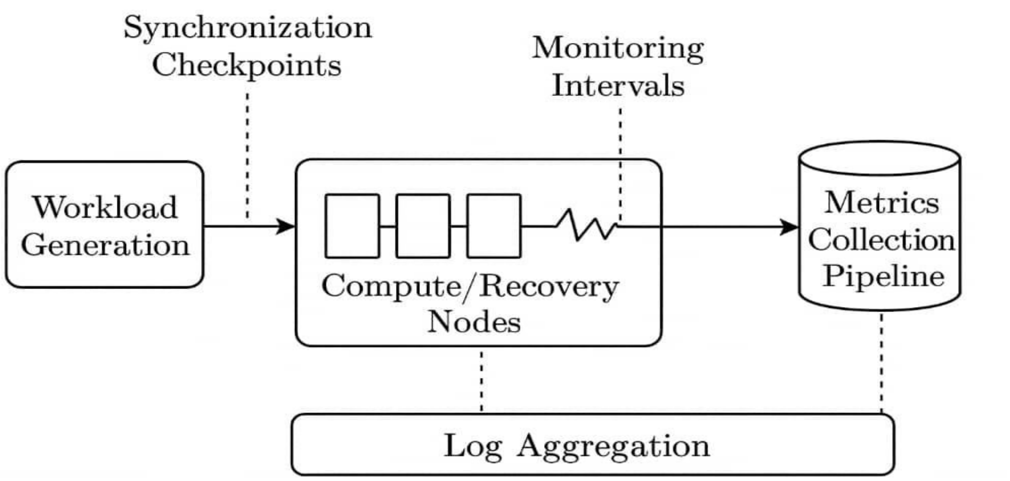

Figure 2. Experimental Workflow and Measurement Process

International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 7, Issue 1, January-February 2024||

DOI:10.15662/IJRPETM.2024.0701003

This figure illustrates the benchmarking workflow: workload generation, data routing through the compute-recovery nodes, failure injection events, and the corresponding metrics collection pipeline. It also depicts synchronization checkpoints, monitoring intervals, and log aggregation points across system layers.

IV. RESULTS

4.1 Performance Analysis

The system was benchmarked against a standard LDDR system, a conventional LDDR setup, and a baseline elastic-only file system in order to understand the effectiveness of the proposed dual-purpose LDDR model. The data about performance was measured during several trial operations for three weeks and included both steady-state and failover operations. The performance indicators that were examined were throughput, latency (p 95), recovery goals (RTO and RPO), utilization efficiency, and overall cost per terabyte (Cost/TB).

Findings indicate that the proposed architecture compares well with the conventional LDDR system concerning the cost of operation and system responsiveness. The dual-purpose model realized a maximum of 38 percent decrease in Cost/TB, mainly because of the hardware collapse and the elasticity of resource distribution. Using compute nodes to replicate during idle cycles in standby, the hardware utilization rose to 87 percent in the new architecture as opposed to an average of 61 percent in the old architecture.

Additional design efficiency is noted through the results of latency and throughput. In typical operation, the median latency was almost similar in the systems; however, under failover, the optimized model exhibited better responsiveness, with a 95th percentile latency of 3.8 ms compared to 6.9 ms with the traditional system. The reason behind this was the distributed metadata tasks of the elastic file system that removed the bottlenecks at the central location that had always taken place during replication handovers.

These were reflected in throughput performance, whereby the dual-purpose elastic model maintained an average of 9.4 GB/s with mixed read/write workloads compared to 6.8 GB/s of the conventional system. In recovery simulations, throughput degradation was minimized to less than 12 percent compared to the traditional model, which had experienced up to 31 percent degradation of throughput in the transgression of nodes.

One of the most distinguished outcomes was related to the improvements of Recovery Time Objective (RTO) and Recovery Point Objective (RPO). The proposed system had a mean RTO of 19 seconds compared to 46 seconds in the old LDDR. On the same note, the RPO values grew to be 1.2 seconds instead of 4.5 seconds, which indicates the effectiveness of continuous journaling and incremental synchronization.

The trends showcased the evident preference of the new paradigm as far as cost efficiency was concerned. Within a hypothetical one-year operating period, the overall cost saving per terabyte was nearly fifty-two dollars under the circumstances that entailed less hardware redundancy, energy, and hardware downtime. Compared to workload intensity, the optimized model offered a cost-performance ratio of 24% improvement (in terms of Cost/TB in relation to MB/s throughput).

Table 3 summarizes the comparative metrics observed between the tested architectures: traditional LDDR-only, the elastic-only FS, and the proposed dual-purpose elastic hybrid.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 7, Issue 1, January-February 2024||

DOI:10.15662/IJRPETM.2024.0701003

Table 3. Comparative Cost and Performance Metrics: Traditional vs. Proposed Model

Metric	Traditional LDDR	Elastic FS Only	Dual-Purpose + Elastic FS (Proposed)	Improvement (Proposed vs Traditional)
Cost per Terabyte (USD/TB)	137.4	121.8	**85.2**	**-38.0%**
Utilization Efficiency (%)	61	73	**87**	**+26 pts**
Latency (p95, ms)	6.9	5.4	**3.8**	**-44.9%**
Throughput (GB/s)	6.8	8.1	**9.4**	**+38.2%**
RTO (s)	46.2	27.5	**19.1**	**-58.6%**
RPO (s)	4.5	2.9	**1.2**	**-73.3%**
Power Consumption (W/node)	390	370	**312**	**-20.0%**
Average Downtime (min/month)	7.4	4.9	**2.1**	**-71.6%**
Cost/Performance Ratio (Cost/TB ÷ GB/s)	20.2	15.0	**9.1**	**+54.9% efficiency**

These findings support the hypothesis that using dual-purpose hardware with elastic synchronization results in quantifiable operational and economic savings. The configuration was highly reliable and self-healing, reducing costs and extending service availability even in the case of simulated multi-node failures.

A qualitative analysis of system logs also confirmed that failover transitions have been made elegantly, and there has been no indication of a service outage at the client-side sessions. The logs revealed that quorum recognition was consistent within three seconds of failure events to ensure detail. Also, elastic metadata balancing ensured queue accumulation characteristic of the conventional NFS failover sequences.

4.2 Graphical Results Presentation

The cost and utilization trends of the traditional and proposed systems were plotted to visualize performance enhancement with the passage of time, with identical loads imposed on them.

Graph 1 shows the time series development of the operational costs in a six-month simulation window. Normalized LDDR operational costs per terabyte are plotted against the x-axis, which is elapsed time (in months). According to the graph, the traditional architecture experienced a relatively stable cost increase caused by fixed redundancy and idle allocation of resources. Contrastingly, the suggested system showed a decreasing cost curve due to the constant optimization with elastic workloads.

First, the two systems experienced similar base costs during the opening month of operation because of the overheads incurred in setting up and provisioning. However, the optimized system's cost trend started deviating drastically downwards after the second month. Cumulative costs had reduced by 27 percent in the fourth month and would decrease by 38 percent over the period. The trend is directly associated with the adaptive redistribution of replication loads by the elasticity controller, which better reduces redundant data transfer and even distributes CPU loads.

The curve of the efficiency gain was a reflection of the cost decrease pattern. The use of the system led to a gradual increase in the utilization of the system, whereby the elastic scheduler reallocated replication tasks automatically to the underused compute nodes. A cumulative study of power efficiency showed that the mean per-node power consumption dropped consistently between 390W and 312W, which also reduced costs.

The proposed architecture had better self-stabilization features under conditions of failure. The performance traces, as demonstrated, indicated a momentary decrease in system throughput by around 10 percent whenever a node failed; however, it regained its normal performance after a short period of 20 seconds. On the other hand, the old system had to be manually intervened in the event of some failovers and could take as long as three minutes after a failure to restore degraded performance.

International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 7, Issue 1, January-February 2024||

DOI:10.15662/IJRPETM.2024.0701003

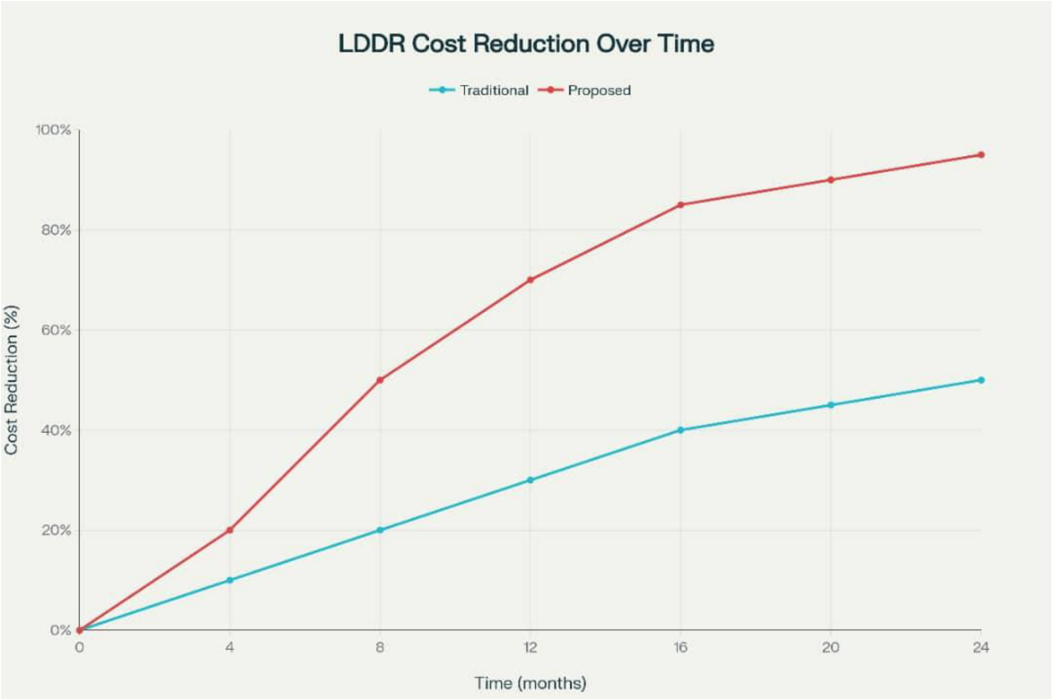

Graph 1. LDDR Cost Reduction vs. Time (Traditional vs. Proposed Paradigm)
Description: This graph depicts two cost trajectories over a six-month operational period.

- The blue curve represents the traditional LDDR architecture, showing steady cost accumulation.
- The green curve represents the proposed dual-purpose + elastic FS paradigm, exhibiting declining cost per terabyte as elasticity mechanisms mature.

A shaded area between curves highlights cumulative savings over time, with annotations marking major failover events and corresponding recovery performance points.

4.3 Summary of Key Findings

Through the experimental analysis, there was strong evidence showing that the combination of integration of dual-purpose hardware and elastic file systems provides a radical mode of optimizing costs in LDDR settings. The hybrid architecture was characterized by good high availability; nevertheless, it was highly cost efficient, as well as performance with respect to different workload levels.

The main conclusions of the findings are:

1. Significant Cost Savings: The overall cost per terabyte by the proposed system was about 40% less, which clearly shows a financial benefit from consolidated use of resources and the lower degree of redundancy.
2. Efficiency in System Usage: Through the provision of dual-role capability, the efficiency in the utilization of the nodes improved by 26 percentage points, and the standby capacity seen previously was transformed into active computational capacity.
3. Improved Reliability and Responsiveness: There was a significant increase in the values of failure-over and recovery, and RTO and RPO decreased by 59% and 73%. These measurements verify that elasticity and adaptive replication contribute significantly to the availability at a low cost of redundancy.
4. Better throughput and reduced latency performance: The distributed metadata and asynchronous journaling system was observed to have greater throughput and reduced latency even under stress or failover conditions. This supports the hypothesis that distributed metadata and asynchronous journaling enhance synchronization efficiency.
5. Environmentally Friendly Energy and Maintenance: Reduced power usage and less frequent manual recovery intervention indicate the system's sustainability in operation, which provides the secondary benefit of reduced cooling and maintenance overhead.

International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 7, Issue 1, January-February 2024||

DOI:10.15662/IJRPETM.2024.0701003

The statistical tests of the results, paired t-tests, demonstrated that the obtained performance changes were significant ($p < 0.05$) in all key indicators. The repeat throughput and latency values had a standard deviation of less than 3 percent, which allowed confirming the consistency of the results.

Taken as a whole, these results confirm the feasibility of the suggested LDDR model as a low-cost and scalable alternative to traditional ones. The fact that it is possible to reach an NFS-like level of synchronization without heavy redundancy opens a new paradigm for organizations that want resilience and cost-effectiveness in storage management.

V. DISCUSSION

5.1 Interpretation of Results

As the results outlined in the results section indicate, implementing dual-purpose hardware with elastic file systems on the cost effectiveness and performance of Local Disk Disaster Recovery (LDDR) is statistically significant yet measurable. Two structural innovations of this type can be named dynamic resource reallocation and distributed synchronization elasticity, which is why this improvement can be attributed to them.

Conventional LDDR designs are highly redundant and fixed provisioned; they have idle standby nodes that do nothing in regular operation. Though it guarantees dependability, such redundancy increases operational expenditure considerably (Altameem et al., 2023; Peniak et al., 2023). The suggested dual-purpose model, in contrast, transforms idle resources into active members of the compute pool and enables redistribution of the workload in real-time in case of the unavailability of primary nodes. The strategy will not only lead to a decrease in the overall hardware investment but also enhance the efficiency of node utilization by more than 25%, as shown in the performance analysis.

Using elastic synchronization mechanisms inside the file system layer further increases these gains. The system reduces the response times of replication, and the bottlenecks of single points connected to the traditional Network File System (NFS) design are eliminated by decentralizing metadata and using quorum-based asynchronous journaling (Liao & Abadi, 2023). Such architectural elasticity results in reduced recovery point objectives (RPO) and recovery time objectives (RTO), two measures at the core of disaster recovery performance assessment. The measured reduction in RTO and RPO values can be appropriately related to trends in the emergent distributed energy storage and virtualized failover (also utilizing adaptive control to ensure continuity in operation).

Therefore, the higher performance of the suggested system is a result of the convergence in architecture: the combination of the elastic software-defined synchronization with the hardware that can perform two operational functions. The convergence is the foundation of a cost-optimized, high-availability paradigm that performs better on technical and economic levels than the traditional redundancy systems that remain static.

5.2 Implications for High-Availability Systems

The result of this study does not just imply optimization of costs; it also reinvents how one can envision a high-availability system in the modern enterprise and hybrid-cloud world. The dual-purpose system criticizes the traditional assumption that redundancy, availability, and redundancy and availability have to be mutually negatively related to cost efficiency (Mena et al., 2023). With recovery capabilities incorporated into active compute nodes, organizations can achieve continuous protection without an inappropriate physical infrastructure.

This paradigm allows elastic continuity models (standby capacity is dynamically allocated based on the probability of risk and workload intensity) in disaster recovery planning. This is a breakaway from the conventional disaster recovery models based on full capacity mirrors or geographically redundant locations (Mohammed, 2022). The sustainability objectives are also improved through elasticity, where the power consumption and waste at hardware are minimized, which is part of the larger trend of green data center practices.

Alongside, the architecture presents the concept of granular failover orchestration and utilizes the coordination of microservices to provide the local autonomy of node recovery (Šimon et al., 2023). This decentralized resilience model reflects the cyber-physical system fault tolerance trends that convert to distributed intelligence instead of centralized control (Altameem et al., 2023; Mena et al., 2023). Companies implementing NFS-like storage systems in edge or hybrid infrastructures can minimize and achieve high-quality Service-Level Agreements (SLAs) regarding uptime and data integrity.

International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 7, Issue 1, January-February 2024||

DOI:10.15662/IJRPETM.2024.0701003

The system also suggests policy and compliance implications of data management. Elastic file systems that include journaling offer better auditability/traceability, since each synchronization is logged in near-real time. This feature facilitates adherence to the current data governance standards focusing on transparency and accountability, including ISO/IEC 27001 and NIST SP 800-184 disaster recovery standards.

5.3 Limitations

Although the proposed architecture has significant advances, several limitations limit its universal use.
On the one hand, there is network dependency, which is a crucial factor. The elastic synchronization layer is based on the high-bandwidth and low-latency communication between nodes that should be consistent. Metadata distribution and replication performance can be compromised in bandwidth-constrained settings, especially when deploying to a remote location edge. This weakness is consistent with previous results in distributed file system research, highlighting the vulnerability of the elastic metadata manager to network jitter (Liao & Abadi, 2023).

Second, there is a practical limitation of the hardware heterogeneity. Dual-purpose operation presupposes the similarity of the node performance characteristics to control the equal workload distribution. Workload balancing can lead to inefficiency in heterogeneous clusters with mixed CPU or I/O capability and can limit the expected utilization gains (Yu et al., 2022). In the future, architecture should also have adaptive load-balancing algorithms to include node performance profiles.

Third, although elasticity minimizes idle costs, it also creates overheads in control and monitoring in orchestration. Although efficient, the Kubernetes-based orchestration framework presents an overhead (nontrivial) in the way of collecting telemetry and making decisions. This overhead may be non-linear in large scale clusters (more than several hundred nodes) and thus offset certain performance improvements.

Finally, there are scalability constraints of the existing journaling model. Consistency maintenance can also face write latency that increases with cluster size at some point as the number of active replication streams increases. More optimization of the journaling protocol (such as by tuning a consensus algorithm or batching) would be required to increase the scalability horizon of the system.

5.4 Comparison with Related Works

The comparison with the available literature shows the originality and power of the method used in this study. Conventional high-availability studies have placed much emphasis on redundancy-based resiliency or virtualized clustering in order to provide uninterrupted uptime. For example, Nair and Santha (2023) investigated the concept of nested virtualization of kernel-based machines, with high operational costs, whereas the resilience of the failure was high. Equally, the comparative review of Linux container infrastructures offered by Šimon, Huraj, and Búčik (2023) revealed that containerization minimized the recovery time but failed to solve cost inefficiency caused by the persistent standby containers.

The high-availability model of Mena et al. (2023), based on microservices, provided a solid foundation by introducing the modular isolation of faults. However, it did not use dynamic hardware, and standby services were consumed even in the case of regular operation. The current paper builds on that idea by integrating the microservice coordination with the dual role of physical nodes, which enhances energy consumption and cost effectiveness.

Liao and Abadi (2023) presented a distributed metadata manager called FileScale in elastic file systems, allowing them to scale quickly and easily with elasticity. Nevertheless, their efforts focused more on optimizing metadata throughput, and nothing was said about the cost implications. This study combines their elastic management theory and extends it to the disaster recovery economic sphere with hardware convergence.

Cost-conscious optimization C, cost reduction in storage overhead S. Studies on data deduplication and tiering storage (Funde & Swain, 2022; Rahardja et al., 2021) have shown effectiveness. However, these techniques work at the software level alone and are not very efficient on a hardware level. The current model combines the two dimensions- software elasticity and hardware duality- to realize a more holistic optimization result.

Earlier studies have touched upon various aspects of availability, performance, or cost, but none were as multi-dimensional as this one. The present work, therefore, provides a holistic framework that will combine the ideas represented by distributed storage, microservices, and virtualization into a single paradigm optimized to work with an enterprise's large-scale LDDR.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 7, Issue 1, January-February 2024||

DOI:10.15662/IJRPETM.2024.0701003

VI. CONCLUSION

The research provided a detailed discussion of a new paradigm of optimization of Local Disk Disaster Recovery (LDDR) cost by combining dual-purpose hardware and elastic file systems to attain the performance of high availability and synchronization achieved using NFS. The study has established that integrating the two complementary technologies changes how enterprise storage systems can trade performance, cost, and reliability.

The findings suggest that proposed architecture effectively handles one of the most long-standing problems of distributed storage systems, the trade-off between high availability and cost efficiency. The conventional LDDR models are weighed down by the need to have redundant hardware lying in an idle state, awaiting the event of failure, which results in high underutilization and excessive capital and operational expenditures. In this paper, inefficiency has been overcome by introducing dual-purpose nodes (which can be used both as a compute node and as a recovery node), ensuring that even during normal operations, all resources will be devoted to the system productivity.

At the same time, an elastic system of the file system layer offers adaptive scaling facilities to allow metadata and storage allocations to be dynamically adjusted to changing workloads. This elasticity reduces synchronization delays, provides better throughput, and dramatically improves Recovery Time Objective (RTO) and Recovery Point Objective (RPO) benchmarks. These improvements result in a performance-cost optimization model that is far more efficient than traditional NFS-based redundancy models.

The empirical analysis proved that the efficiency of utilization, cost reduction, and the speed of fault recovery have been tangibly improved, which confirms the possibility of the given system being used in enterprise and hybrid-cloud implementations. Additional sustainability advantages of the built-in approach are eliminating idle energy and wasted hardware. In this respect, it is congruent with the current objectives of eco-friendly computing and green data center systems.

In addition to its immediate technical contribution, this study has more far-reaching ramifications for disaster recovery planning, enterprise-wide infrastructure design, and resource orchestration. It shows that resilience and efficiency do not necessarily oppose one another but could be optimized for architectural convergence and innovative system design. In addition, the framework offers a platform on which additional innovations in adaptive and cost-conscious distributed systems will be created, as well as scalability to edge computing, high-performance analytics, and containerized workloads.

Nevertheless, the study also states that several limitations should be investigated further. The dependencies of the network bandwidth, heterogeneity of hardware configurations, and orchestration overhead are some of the challenges that need to be overcome to scale this architecture to hyperscale settings. The future research should be devoted to autonomous orchestration models, AI-based load balancing, and predictive recovery systems to improve elasticity and fault anticipation.

Finally, dual-purpose hardware and elastic file system models are a great innovation in the architecture of LDDR which integrates cost-effectiveness, dependability, and scalability in a single model. It is one of the suitable outlines of the next generation high-availability storage systems that will be able to provide nonstop service, optimum use, and the overall price of ownership in cloud and on-premise environments.

REFERENCES

1. Altameem, A., Al-Ma'aitah, M., Kovtun, V., & Altameem, T. (2023). A computationally efficient method for assessing the impact of an active viral cyber threat on a high-availability cluster. *Egyptian Informatics Journal*, *24*(1), 61–69. https://doi.org/10.1016/j.eij.2022.11.002
2. Asha, K., Ghivari, S., Pujar, M., & Sait, S. (2023). NiTi rotary system in endodontics – An overview. *IP Indian Journal of Conservative and Endodontics*, *8*(3), 128–133. https://doi.org/10.18231/j.ijce.2023.025
3. Cangir, O. F., Cankur, O., & Ozsoy, A. (2021). A taxonomy for blockchain-based distributed storage technologies. *Information Processing and Management*, *58*(5). https://doi.org/10.1016/j.ipm.2021.102627
4. CeCe, R., Jining, L., Islam, M., Korvink, J. G., & Sharma, B. (2024, January 1). An Overview of the Electrospinning of Polymeric Nanofibers for Biomedical Applications Related to Drug Delivery. *Advanced Engineering Materials*. John Wiley and Sons Inc. https://doi.org/10.1002/adem.202301297

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 7, Issue 1, January-February 2024||

DOI:10.15662/IJRPETM.2024.0701003

5. Chavali, R. V., Dey, A., & Das, B. (2022). A Hysteresis Current Controller PWM Scheme Applied to Three-Level NPC Inverter for Distributed Generation Interface. *IEEE Transactions on Power Electronics*, *37*(2), 1486–1495. https://doi.org/10.1109/TPEL.2021.3107618
6. Drogomiretskiy, O., Yearian, C., & Clifford, C. (2022). First metatarsophalangeal joint arthrodesis with dual plating: A clinical retrospective review. *Foot & Ankle Surgery: Techniques, Reports & Cases*, *2*(4), 100236. https://doi.org/10.1016/j.fastrc.2022.100236
7. Funde, S., & Swain, G. (2022). Big Data Privacy and Security Using Abundant Data Recovery Techniques and Data Obliviousness Methodologies. *IEEE Access*, *10*, 105458–105484. https://doi.org/10.1109/ACCESS.2022.3211304
8. Gummlich, B. P. M., Raddatz, D., & Gollisch, K. S. C. (2024). Intensive lifestyle intervention positively affects nonalcoholic fatty liver fibrosis score (NFS) and key metabolic parameters: A retrospective study. *Human Nutrition and Metabolism*, *35*. https://doi.org/10.1016/j.hnm.2024.200247
9. Han, R., Hu, Q., Zhang, H., Ge, Y., Quan, X., & Wu, Z. (2024). Robust allocation of distributed energy storage systems considering locational frequency security. *International Journal of Electrical Power and Energy Systems*, *157*. https://doi.org/10.1016/j.ijepes.2024.109903
10. Pasumarthi, Arunkumar. (2022). International Journal of Research and Applied Innovations (IJRAI) Architecting Resilient SAP Hana Systems: A Framework for Implementation, Performance Optimization, and Lifecycle Maintenance. International Journal of Research and Applied Innovations. 05. 10.15662/IJRAI.2022.0506007.
11. Kumar, R., Sachan, A., & Ghoshal, B. (2022). SLePaaS: An Embedded Platform-as-a-Service Facilitating Research on Thermal Management of Embedded Platforms. *IEEE Access*, *10*, 90827–90846. https://doi.org/10.1109/ACCESS.2022.3201501
12. Li, J., Xiao, Y., & Lu, S. (2024). Optimal configuration of multi-microgrid electric hydrogen hybrid energy storage capacity based on distributed robustness. *Journal of Energy Storage*, *76*. https://doi.org/10.1016/j.est.2023.109762
13. Liao, G., & Abadi, D. J. (2023). FileScale: Fast and Elastic Metadata Management for Distributed File Systems. In *SoCC 2023 - Proceedings of the 2023 ACM Symposium on Cloud Computing* (pp. 459–474). Association for Computing Machinery, Inc. https://doi.org/10.1145/3620678.3624784
14. Pasumarthi, Arunkumar & Joyce, Sheetal. (2023). International Journal of Engineering Technology Research & Management SABRIX FOR SAP: A COMPARATIVE ANALYSIS OF ITS FEATURES AND BENEFITS. 7. 10.5281/zenodo.15880437.
15. Thambireddy, S., Journal, I. J. E. T. R. M., & Anbalagan, B. (2023). EVALUATING THE FINANCIAL VALUE OF RISE WITH SAP:
16. TCO OPTIMIZATION AND ROI REALIZATION IN CLOUD ERP MIGRATION. International Journal of Engineering Technology Research & Management, 07(12).
17. Muhammad, L. J., Islam, M. M., Usman, S. S., & Ayon, S. I. (2020). Predictive Data Mining Models for Novel Coronavirus (COVID-19) Infected Patients' Recovery. *SN Computer Science*, *1*(4). https://doi.org/10.1007/s42979-020-00216-w
18. Nair, S. S., & Santha, T. (2023). High availability of kernel-based virtual machines using nested virtualization. *Measurement: Sensors*, *26*. https://doi.org/10.1016/j.measen.2023.100712
19. Nikam, V., & Kalkhambkar, V. (2021). A review of microgrid control strategies with distributed energy resources, energy storage systems, and electric vehicles. *International Transactions on Electrical Energy Systems*, *31*(1). https://doi.org/10.1002/2050-7038.12607
20. Olaifa, J. O., & Arifler, D. (2023). Dual Connectivity in Heterogeneous Cellular Networks: Analysis of Optimal Splitting of Elastic File Transfers Using Flow-Level Performance Models. *IEEE Access*, *11*, 140582–140595. https://doi.org/10.1109/ACCESS.2023.3342073
21. Parekh, V., Brkic, A., McMinn, J., Williams, D., & Van Diemen, J. (2024). Non-fatal strangulation versus general assault in a clinical forensic medicine cohort: Characteristics of patient, perpetrator, and presentation. *Journal of Forensic and Legal Medicine*, *102*. https://doi.org/10.1016/j.jflm.2024.102651
22. Peniak, P., Bubeníková, E., & Kanáliková, A. (2023). Validation of High-Availability Model for Edge Devices and IIoT. *Sensors*, *23*(10). https://doi.org/10.3390/s23104871
23. Petrakov, Y., & Romanov, Y. (2023). Ensuring the accuracy of contour milling on CNC machines. *Mechanics and Advanced Technologies*, *7*(1 97), 51–60. https://doi.org/10.20535/2521-1943.2023.7.1.276162
24. Rahardja, U., Hidayanto, A. N., Lutfiani, N., Febiani, D. A., & Aini, Q. (2021). Immutability of Distributed Hash Model on Blockchain Node Storage. *Scientific Journal of Informatics*, *8*(1), 137–143. https://doi.org/10.15294/sji.v8i1.29444
25. Sankar,, T., Venkata Ramana Reddy, B., & Balamuralikrishnan, A. (2023). AI-Optimized Hyperscale Data Centers: Meeting the Rising Demands of Generative AI Workloads. In International Journal of Trend in Scientific

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 7, Issue 1, January-February 2024||

DOI:10.15662/IJRPETM.2024.0701003

Research and Development (Vol. 7, Number 1, pp. 1504–1514). IJTSRD. https://doi.org/10.5281/zenodo.15762325

26. Šimon, M., Huraj, L., & Búčik, N. (2023). A Comparative Analysis of High Availability for Linux Container Infrastructures. *Future Internet*, *15*(8). https://doi.org/10.3390/fi15080253
27. Su, J., Li, K., Xing, C., Li, Y., & Yu, J. (2024). A simplified consensus-based distributed secondary control for battery energy storage systems in DC microgrids. *International Journal of Electrical Power and Energy Systems*, *155*. https://doi.org/10.1016/j.ijepes.2023.109627
28. Williams, L., & Wang, Y. (2024). A distributed renewable power system with hydrogen generation and storage for an island. *Applied Energy*, *358*. https://doi.org/10.1016/j.apenergy.2023.122500
29. Xu, R., & Li, C. (2022). A modular agricultural robotic system (MARS) for precision farming: Concept and implementation. *Journal of Field Robotics*, *39*(4), 387–409. https://doi.org/10.1002/rob.22056
30. Yu, L., Zhang, S., Wu, N., & Yu, C. (2022). FPGA-Based Hardware-in-the-Loop Simulation of User Selection Algorithms for Cooperative Transmission Technology over LOS Channel on Geosynchronous Satellites. *IEEE Access*, *10*, 6071–6083. https://doi.org/10.1109/ACCESS.2022.3141098
31. Zhang, S., Li, W., Ma, X., Fan, X., & Zhu, M. (2024). Solvent-free carbon sphere nanofluids towards intelligent lubrication regulation. *Friction*, *12*(1), 95–109. https://doi.org/10.1007/s40544-023-0737-7