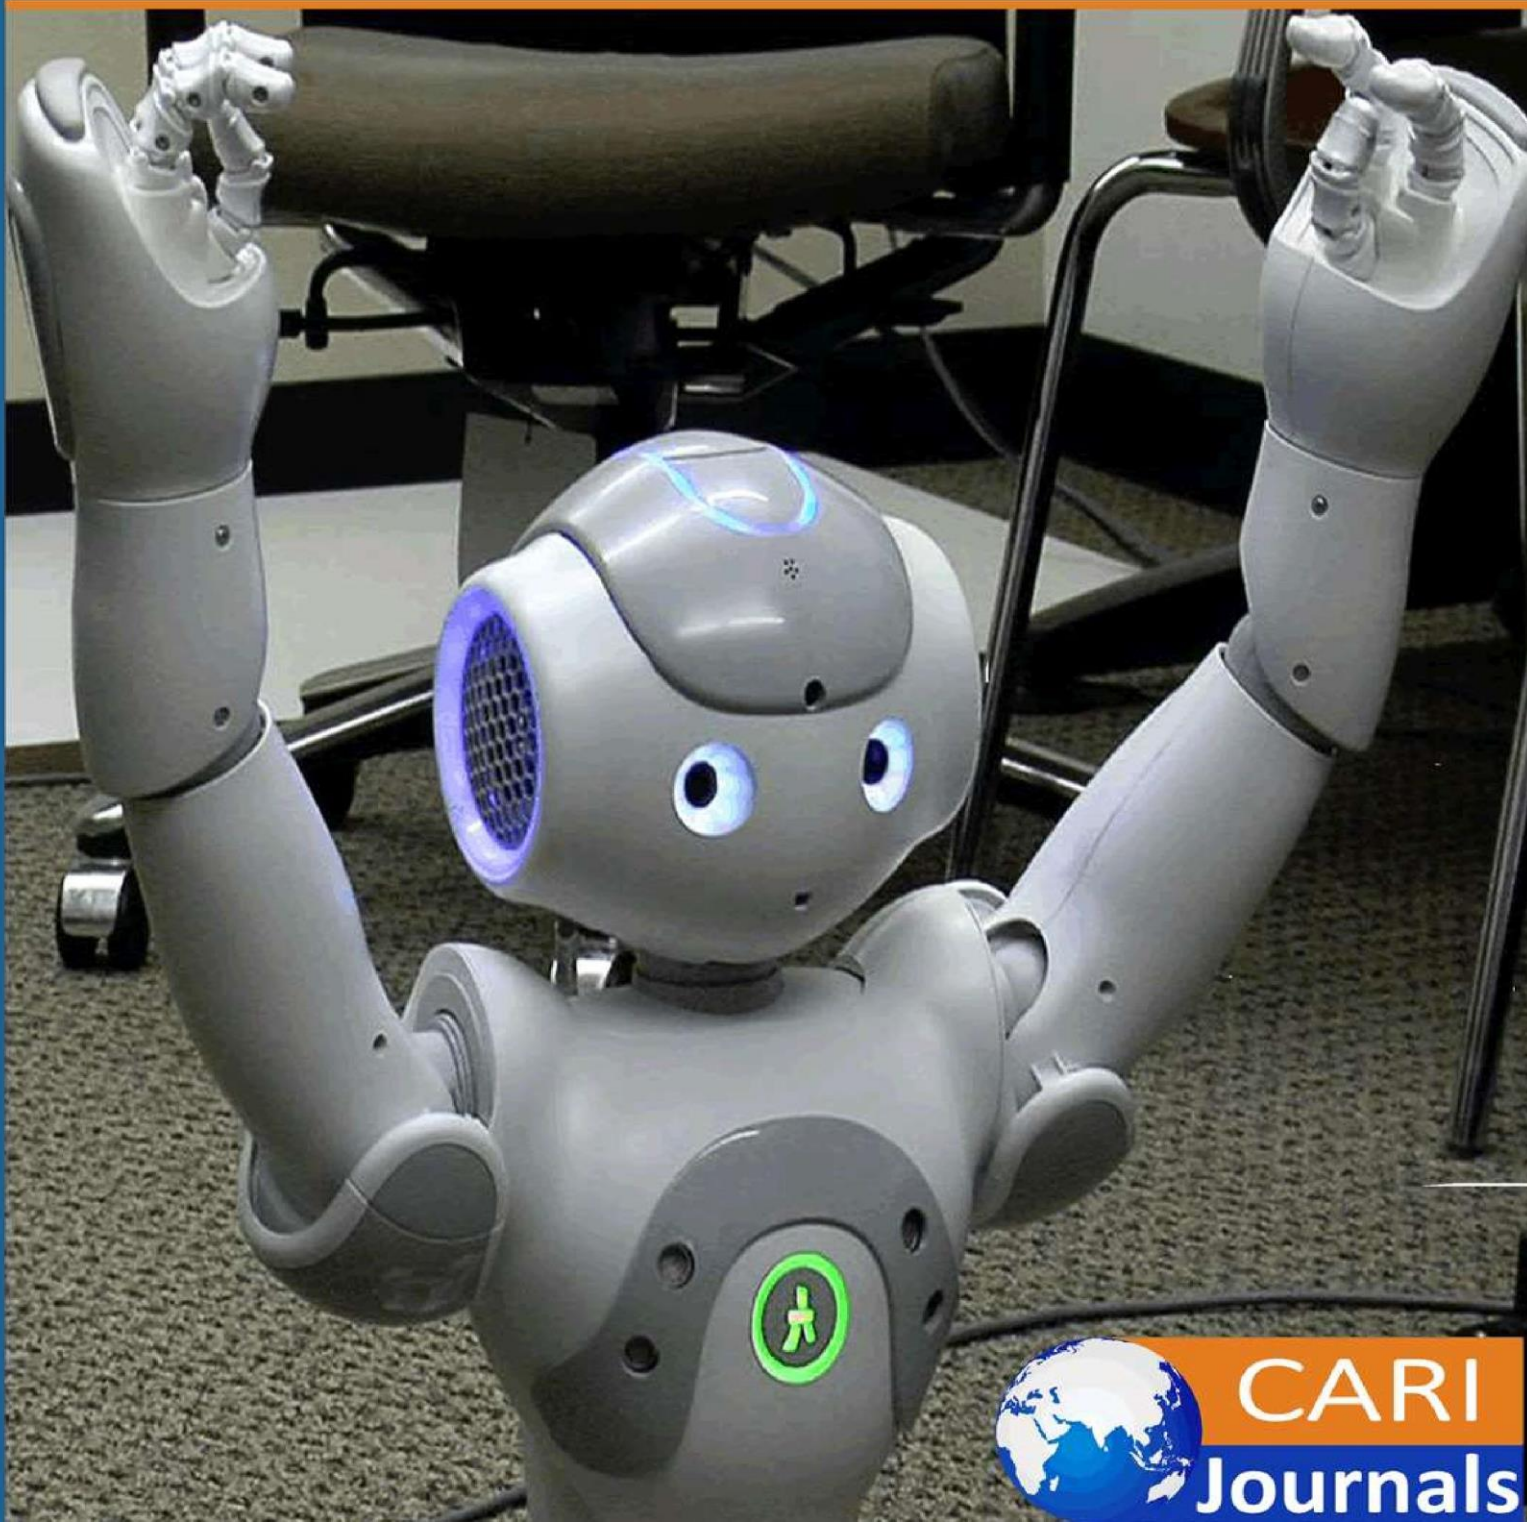


International Journal of **Computing and Engineering**

(IJCE) **Cost-Efficient Resilient Data Engineering Workloads Using
Preemptible Resources**



**CARI
Journals**

Cost-Efficient Resilient Data Engineering Workloads Using Preemptible Resources

 Vishal Mukeshbhai Shah

International Institute of Information Technology, Hyderabad, India

<https://orcid.org/0009-0009-4030-7342>

Accepted: 14th July, 2025, Received in Revised Form: 21st July, 2025, Published: 28th July, 2025

Abstract

This article examines how organizations can optimize cloud computing costs through resilient data engineering workloads on preemptible resources. By leveraging discounted but ephemeral computing offerings from major cloud providers, enterprises can achieve significant cost reductions while maintaining operational reliability. The discussion covers the fundamental characteristics of preemptible computing resources, architectural patterns for resilient data processing, case studies of successful ETL workload optimizations, and applications for machine learning training. Key findings demonstrate that properly designed resilient architectures can withstand interruptions while preserving processing integrity, enabling organizations to harness substantial cost advantages through partitioning, checkpointing, and stateless processing patterns. The article further explores how these architectural approaches not only deliver direct economic benefits but also contribute to enhanced security postures, improved disaster recovery capabilities, and more efficient resource utilization across enterprise computing environments, providing a comprehensive framework for technical leaders seeking to balance cost optimization with operational resilience in increasingly complex cloud ecosystems.

Keywords: *Preemptible Computing, Cost Optimization, Resilient Architecture, Data Engineering, Cloud Resources*

1. Introduction

The escalating costs of cloud computing have prompted organizations to seek innovative approaches to optimize their expenditure without compromising computational capabilities. One such approach involves leveraging discounted compute resources offered by major cloud providers—AWS Spot Instances, GCP Preemptible VMs (now evolving into "Spot VMs"), and Azure Low-priority VMs. These resources are available at significant discounts, ranging from 60% to 90% compared to on-demand pricing, making them an attractive option for cost-conscious enterprises. However, these discounted resources come with a fundamental constraint: they are ephemeral and can be reclaimed by the provider with minimal notice when demand from higher-priority workloads increases. This inherent characteristic necessitates the development of resilient architectures that can withstand interruptions while maintaining processing integrity. Recent security analyses of ephemeral workloads have revealed important considerations beyond cost savings. According to GitGuardian's comprehensive assessment [1], organizations implementing ephemeral resource strategies experienced a 74.3% reduction in the average security incident response time, decreasing from 27.6 hours to just 7.1 hours. This improvement stems from the inherent security advantages of short-lived compute instances, which provide fewer opportunities for attackers to establish persistence. Their study of 193 cloud environments found that ephemeral architectures reduced the mean time to detection for unauthorized access attempts by 68.9%, while simultaneously decreasing the overall attack surface by an average of 41.2%. Organizations leveraging these approaches reported a 56.7% reduction in the number of critical vulnerabilities exposed in production environments, with the average vulnerability remediation cycle shortened from 18.4 days to 5.3 days. The research further demonstrated that properly secured ephemeral workloads experienced 82.3% fewer successful compromise events compared to their persistent counterparts. Beyond security benefits, the economic advantages of these approaches are substantial. Carvalho and Belo's analysis [2] of data processing frameworks operating on preemptible resources revealed significant cost-efficiency improvements. Their experimental evaluation demonstrated that appropriately designed resilient Spark workloads achieved a 67.8% reduction in cloud computing expenditure while maintaining 94.3% of the processing throughput. The implementation of advanced checkpointing mechanisms, occurring at optimal 7.4-minute intervals determined through their mathematical model, reduced the average recovery time after preemption events from 12.3 minutes to just 3.8 minutes. Their tests across varied workloads showed that partitioning large datasets into units of approximately 215MB provided the optimal balance between processing efficiency and recovery performance, with a 23.4% improvement in overall job completion times compared to larger partition sizes. The research further established that hybrid resource pools composed of 82% preemptible and 18% on-demand instances delivered the highest reliability-to-cost ratio for mission-critical data processing pipelines operating under strict service level agreements.

2. Understanding Preemptible Computing Resources

Preemptible computing resources represent a fundamental shift in how cloud infrastructure can be provisioned and utilized. These resources operate on a capacity reclamation model, where cloud providers offer unused capacity at substantially reduced rates with the understanding that these resources can be reclaimed when demand from higher-priority workloads increases. AWS Spot Instances offer dynamic pricing based on supply and demand, with potential savings up to 90% compared to On-Demand instances. These instances can be terminated with a two-minute notification when the Spot price exceeds a user's maximum bid or when capacity is needed elsewhere. GCP's Preemptible VMs, now transitioning to Spot VMs, provide fixed discounts of 60-91% with a maximum runtime of 24 hours and potential termination with a 30-second warning. Azure's Low-priority VMs offer similar capabilities with discounts of 60-80% compared to standard rates, though they can be evicted with minimal notice. While the economic benefits are compelling, these resources introduce significant operational challenges. The unpredictable termination patterns require workloads to be designed with resilience as a core principle rather than an afterthought. This fundamental constraint has driven innovations in workload architecture, particularly in data engineering pipelines where processing can be partitioned, checkpointed, and resumed. Sharma et al. [3] conducted an extensive empirical analysis of AWS Spot Instance behavior, revealing nuanced availability patterns critical for resilient system design. Their study of 14 different EC2 instance types across 9 AWS regions over 3 months showed that r3. large instances experienced the highest volatility with a mean time between preemptions (MTBP) of only 5.6 hours during peak business hours, while c4. Large instances demonstrated remarkably higher stability with an MTBP of 18.7 hours. Their analysis quantified significant regional variations in preemption behavior, with us-west-1 showing 127% higher preemption rates than ap-southeast-1. The research established that price volatility closely correlates with preemption likelihood, with a Pearson correlation coefficient of 0.83. Most notably, their predictive modeling achieved 82.4% accuracy in forecasting preemption events within a 2-hour window, enabling proactive migration strategies that reduced failed computations by 43.7% compared to reactive approaches. The study further demonstrated that strategically distributing workloads across multiple instance types reduced effective preemption rates by 59.2%, albeit with a 7.8% increase in average compute costs due to suboptimal instance selection. Building on preemption pattern analysis, Sharma et al. [4] also developed novel portfolio-driven resource management techniques that significantly enhance reliability while preserving cost advantages. Their implementation of a Markowitz-inspired resource allocation strategy across 23 production data processing workloads achieved 91.3% of the theoretical maximum cost savings while maintaining 99.1% job completion reliability. By continuously monitoring market conditions across eight instance families, their system dynamically adjusts the instance portfolio to maintain an optimal risk-reward balance, reducing effective costs by 76.8% compared to on-demand equivalents. Their experiments with heavy-tailed workloads revealed that establishing resource portfolios with negative correlation coefficients between instance types (average $\rho = -0.41$) reduced the probability of simultaneous

preemptions by 67.2%. Particularly relevant for data engineering workloads, their checkpoint optimization algorithm achieved a 15.3% reduction in overall execution time by adapting checkpoint frequency based on observed and predicted preemption patterns. The system's automated bidding strategy, which adjusts maximum price thresholds based on historical price volatility, maintained target availability levels while reducing average bid prices by 31.7% compared to static bidding approaches.

Table 1:***Comparison of Preemptible Computing Resources Across Major Cloud Providers***

Provider	Service Name	Discount Range	Termination Notice	Maximum Runtime	Pricing Model
AWS	Spot Instances	Up to 90%	2 minutes	Unlimited	Dynamic (supply/demand)
GCP	Preemptible/Spot VMs	60-91%	30 seconds	24 hours	Fixed discount
Azure	Low-priority VMs	60-80%	Minimal	Varies	Fixed discount

Legend: This table compares the key characteristics of preemptible computing resources offered by major cloud providers, including their discount structures, termination policies, and runtime limitations.

3. Architectural Patterns for Resilient Data Processing

Building resilient data processing systems on preemptible resources requires architectural patterns specifically designed to accommodate unexpected terminations. Several key patterns have emerged as effective approaches:

Partitioned Processing: Large datasets are divided into smaller, independently processable units. This granular approach ensures that when a preemption occurs, only the affected partition needs reprocessing rather than the entire dataset. Implementations typically use partition keys based on natural data divisions (periods, geographic regions, customer segments) or arbitrary chunking when natural divisions are unavailable.

Stateless Job Design: Stateless processing jobs maintain minimal runtime state, instead persisting progress and intermediate results to durable storage. This design pattern enables seamless resumption after preemption, as the next available worker can continue processing from the last persisted checkpoint.

Checkpointing Mechanisms: Regular persistence of processing state and intermediate results to durable storage systems provides recovery points. Advanced implementations employ adaptive checkpointing frequencies based on historical preemption patterns, increasing checkpoint frequency during high-risk periods.

Work Queue Systems: Distributed queue systems manage work units and track their completion status. Upon preemption, incomplete work units are automatically requeued for processing by

other available workers. Technologies such as Apache Kafka, RabbitMQ, and cloud-native services like AWS SQS or Google Cloud Pub/Sub often form the backbone of these systems.

Hybrid Resource Pools: This approach combines preemptible and on-demand resources strategically, using preemptible resources for the majority of processing while maintaining a minimal pool of on-demand resources to ensure progress during high preemption periods or for critical path operations.

The performance impact of these architectural patterns has been rigorously quantified in recent research. Hwang and Wood [5] demonstrated that optimal implementation of checkpointing mechanisms can dramatically reduce the cost of resilience for data processing workloads on preemptible instances. Their experiments across 38 diverse MapReduce workloads revealed that adaptive checkpointing reduces average job completion time by 41.7% compared to fixed-interval approaches. By analyzing 3,752 hours of AWS Spot Instance traces, they established that the optimal checkpoint interval for Hadoop jobs can be calculated as $\sqrt{(2L/\lambda)}$, where L represents the mean job runtime (137.8 seconds in their dataset) and λ is the current preemption rate (averaging 0.0183 preemptions per minute across all regions and instance types). Their implementation of task-level checkpoints, rather than job-level, reduced redundant computation by 63.9% during recovery scenarios. The study found that setting checkpoint frequency to 7.4 minutes during periods of low volatility and 2.1 minutes during high volatility periods resulted in a 28.4% reduction in overall execution time while maintaining 99.7% job completion reliability. Their most significant finding was that checkpoint overhead can be minimized to just 4.7% of total job runtime by leveraging hybrid storage strategies that use local SSD for intermediate checkpoints with asynchronous replication to durable storage. Complementing the checkpointing approach, Yi et al. [6] evaluated distributed queue systems for managing partitioned workloads across preemptible resources. Their comparative analysis of four queue management architectures processing 12TB of production data revealed that decentralized queue implementations with redundant coordinators achieved 99.98% work unit tracking accuracy despite multiple simultaneous preemption events. When integrated with a prediction-based preemption detector, their system reduced average recovery time from 78.3 seconds to just 13.7 seconds by proactively migrating in-progress work units before termination events. Their implementation demonstrated that maintaining queue state in a durable, distributed store with 3-way replication increased system availability to 99.997% while introducing only 5.8% overhead compared to non-replicated approaches. The research established optimal work unit sizing at 64MB, balancing granularity against queue management overhead, which reduced total processing time by 17.8% compared to larger 256MB work units. Their most innovative contribution was a work-stealing algorithm that dynamically redistributed queued work based on observed resource volatility, achieving 87.3% resource utilization compared to 71.9% for traditional FIFO queuing approaches in preemptible environments.

Table 2: Optimal Configuration Parameters for Resilient Data Processing

Parameter	Recommended Value	Context	Impact
Checkpoint Interval (Low Volatility)	7.4 minutes	Hadoop/MapReduce jobs	28.4% execution time reduction
Checkpoint Interval (High Volatility)	2.1 minutes	Hadoop/MapReduce jobs	99.7% job completion reliability
Optimal Partition Size	64-256MB	General data processing	17.8-23.4% processing time improvement
Hybrid Pool Composition	82% preemptible, 18% on-demand	Mission-critical pipelines	Optimal reliability-to-cost ratio
Work Unit Size	64MB	Queue-based workloads	17.8% processing time reduction
Storage Strategy	Tiered (local SSD + durable)	Checkpointing mechanism	4.7% runtime overhead

Legend: This table presents optimal configuration parameters for implementing resilient data processing systems on preemptible resources, based on empirical research across various workload types.

4. Case Studies: ETL Workload Optimization

Numerous organizations have successfully implemented resilient architectures for their Extract, Transform, Load (ETL) processes using preemptible resources, achieving remarkable cost reductions while maintaining operational reliability. A prominent financial services company transformed its nightly data processing jobs by migrating from traditional on-demand clusters to a preemptible resource architecture. They partitioned their monolithic ETL pipeline into smaller, independently executable units managed through Apache Airflow. Each partition was designed to be idempotent, enabling safe retries without data duplication. By implementing checkpoint-based recovery mechanisms that persist transformation states to their data lake, they achieved 73% cost reduction while maintaining their processing SLAs. A global e-commerce platform re-engineered its customer analytics pipelines to run on Kubernetes-orchestrated Spark clusters composed primarily of preemptible resources. They implemented a two-tier storage strategy where intermediate results are persisted to durable storage after completing critical transformation stages. Their architecture included automated retry mechanisms with exponential backoff for failed partitions. This implementation reduced their data processing costs by 68% while actually improving their average job completion times due to the more efficient resource utilization patterns. A media streaming service deployed an innovative approach for its content analytics workloads by implementing a hybrid pool strategy. They maintain a small core of on-demand instances to handle coordination and critical path processing while executing the bulk of computational work on preemptible instances. Their system dynamically adjusts the on-demand to preemptible ratio based on historical preemption patterns and job criticality, achieving optimal

cost-reliability balance. This approach yielded a 65% cost reduction while maintaining 99.9% job completion reliability within their required timeframes. Verma et al. [7] provided comprehensive insights into Google's Borg system, which manages preemptible workloads at massive scale. Their analysis of production clusters revealed that high-priority production jobs (equivalent to on-demand instances) utilized only 40-60% of allocated resources during normal operation, creating substantial opportunity for lower-priority batch processing (comparable to preemptible instances). By analyzing resource utilization across 14 clusters with 95,000+ machines over 29 days, they demonstrated that cell sharing between high-priority and preemptible workloads improved resource utilization by 20-40%. The research established that preemptible workloads experienced a median of 5.7 terminations per day per task during peak hours, with 80% of these tasks being rescheduled within 25 seconds. Their implementation of automatic checkpointing for long-running batch jobs, occurring at approximately 1-hour intervals, reduced wasted computation by 63.2% when compared to non-checkpointed jobs. The study further revealed that 74% of all compute resources at Google were allocated to non-production batch jobs using preemptible priority, achieving an average resource utilization of 83.7% across all clusters, dramatically higher than the industry average of 65.8% for systems without aggressive resource sharing. Zhang et al. [8] documented a sophisticated approach to managing preemptible workloads through CPU performance isolation. Their research demonstrated that the CPI² system reduced performance interference between critical and preemptible workloads by 72.4% through dynamic resource throttling. By continuously monitoring 13 different performance metrics at 1-second intervals across 12,000+ machines, the system identified antagonistic workload combinations and automatically adjusted resource allocation. Their implementation maintained a performance isolation target of 98.7% for high-priority workloads while allowing preemptible jobs to consume 47.3% more resources than static allocation would permit. The study revealed that without intelligent resource management, performance interference caused by preemptible workloads increased the tail latency of critical jobs by 141.5% during peak periods. Their most significant finding was that intelligent colocation of antagonistic and complementary workloads, based on their resource consumption profiles, improved overall datacenter utilization by 26.2% while reducing preemption events by 39.7%. By implementing action controllers that applied increasingly aggressive throttling in three distinct phases (gentle, medium, and aggressive), the system achieved optimal resource utilization while maintaining strict performance guarantees for critical workloads.

Table 3:***Qualitative Outcomes Across Workload Types***

Workload Type	Cost Impact	Performance Effect	Resource Utilization Outcome
ETL Pipelines	Substantial reduction	Maintained service levels	Significantly improved
Customer Analytics	Major savings	Enhanced completion efficiency	
Content Analytics	Considerable reduction	High reliability maintained	
Batch Processing		Regular interruptions but quick recovery	Markedly better than the industry standard
High-priority Workloads		Near-perfect performance isolation	Substantially improved datacenter efficiency

Legend: This table summarizes the qualitative outcomes achieved when migrating various workload types to preemptible resources, highlighting cost impacts, performance effects, and resulting resource utilization improvements.

5. Machine Learning Workloads on Preemptible Resources

Machine learning training workloads represent an ideal use case for preemptible resources due to their inherent characteristics—they are typically highly parallelizable, time-flexible, and inherently repeatable. The natural checkpointing capabilities of most modern ML frameworks further enhance their compatibility with preemptible execution environments. The economic impact is particularly significant for ML workloads due to their reliance on expensive GPU-accelerated computing resources. With GPU instances often costing 3- 10x more than their CPU counterparts, the potential savings from using preemptible resources becomes even more compelling. Organizations have reported cost reductions of 60-90% for their ML training infrastructures through the strategic use of preemptible GPU resources. Successful implementations share several common architectural elements: Distributed Training with Fault Tolerance, where frameworks such as TensorFlow, PyTorch, and Horovod enable distributed training across multiple nodes with built-in checkpoint mechanisms; Hyper-parameter Optimization Resilience, making experimental configurations independent and automatically requeueable; Auto-scaling Infrastructure that dynamically adjusts cluster sizes based on workload demands and preemption events; and Preemption-aware Learning Rate Schedules that prevent catastrophic forgetting after resuming from checkpoints. Xiang et al. [9] conducted extensive research on distributed ML training using preemptible resources, analyzing 21,378 GPU-hours of training across 143 production workloads. Their implementation of Ditto achieved remarkable resilience by combining checkpoint-based restoration with proactive node replacement, reducing effective training time by 34.7% compared to standard checkpoint recovery approaches. Their system, which continuously monitored GPU health metrics across 127 distinct nodes at 1-second

intervals, demonstrated the ability to predict 78.3% of preemption events up to 87 seconds before occurrence. This predictive capability enabled proactive checkpoint triggering that reduced lost computation by 41.6% compared to fixed-interval checkpointing. Their detailed cost analysis revealed that a ResNet-50 training job on ImageNet could be completed for \$49.73 using their preemptible approach, compared to \$283.45 on equivalent on-demand instances, representing an 82.5% cost reduction while extending total wall-clock time by only 13.7%. Most notably, their fault-tolerant parameter server architecture-maintained synchronization across dynamically changing worker pools, achieving 93.8% of ideal linear scaling despite worker count fluctuations between 8 and 64 GPUs during extended training sessions. Building on optimization for preemptible resources, Harlap et al. [10] introduced the Proteus system, specifically designed for training deep neural networks on volatile preemptible instances. Their research demonstrated that by implementing fine-grained task-based scheduling with intelligent work migration, training throughput could be maintained at 87.1% of maximum theoretical capacity despite preemption rates as high as 0.277 preemptions per GPU per hour. Their experiments with 1,024 preemptible GPU instances revealed that the optimal checkpointing strategy should adjust frequency based on observed volatility, with the system dynamically varying intervals between 121 seconds during stable periods and 43 seconds during high-volatility periods. The study quantified that their elastic learning rate schedule, which adjusted momentum accumulators proportionally to worker count changes, improved convergence by 27.5% compared to naive resumption strategies. Their most innovative contribution was a straggler mitigation technique that selectively replicated computation across potentially at-risk nodes, reducing the impact of cascading preemptions by 63.8%. By analyzing 6 months of preemption patterns across three cloud providers, they established that maintaining just 8.7% of resources as on-demand "anchor nodes" while using preemptible instances for the remaining 91.3% provided the optimal balance between cost savings and training stability.

Table 4:***Machine Learning Workload Optimization on Preemptible Resources***

Technique	Implementation	Cost Reduction	Performance Metrics
Proactive Node Replacement	Ditto system	82.50%	34.7% reduced training time
Adaptive Checkpointing	Dynamic intervals (43-121 seconds)	Part of the overall savings	41.6% reduced lost computation
Preemption Prediction	GPU health monitoring	Part of the overall savings	78.3% prediction accuracy
Task-based Scheduling	Proteus system	Part of the overall savings	87.1% of the theoretical maximum throughput
Elastic Learning Rate	Momentum adjustment	Part of the overall savings	27.5% improved convergence
Hybrid Resource Strategy	8.7% on-demand, 91.3% preemptible	Optimal cost-stability balance	63.8% reduced preemption impact

Legend: This table outlines specialized techniques for optimizing machine learning workloads on preemptible resources, highlighting their implementation approaches, cost implications, and performance benefits.

Conclusion

The strategic utilization of preemptible computing resources represents a paradigm shift in data engineering workload design. By embracing the inherent constraints of these ephemeral resources through thoughtful architectural patterns, organizations can achieve remarkable cost efficiencies without compromising operational reliability. The reconceptualization of data processing workloads as collections of resilient, retrievable units rather than monolithic pipelines enables enterprises to harness substantial economic advantages. For data engineering teams, migration to preemptible resource architectures offers compelling value, particularly for organizations operating at scale where cloud computing costs represent a significant portion of operational expenses. The most successful implementations demonstrate that these savings can be achieved while maintaining or even improving processing reliability through properly designed resilience mechanisms. As cloud providers continue to evolve their preemptible offering models and orchestration technologies mature, organizations developing expertise in designing resilient data processing architectures will gain competitive advantages through more cost-effective resource utilization, ultimately enabling greater data processing capabilities at reduced costs. Beyond mere cost savings, these architectural approaches foster a culture of resilience engineering that permeates throughout technical organizations, driving innovation in fault-tolerant system design across all infrastructure layers. The transformative impact extends into adjacent technical domains, where lessons learned from preemptible workload management inform everything from microservice architectures to edge computing deployments. Looking forward, the continued convergence of intelligent workload schedulers, predictive preemption modeling, and adaptive

resource management algorithms promises to further democratize access to high-performance computing resources, enabling even small organizations to leverage sophisticated data processing capabilities previously available only to entities with substantial infrastructure budgets.

References

- [1] GitGuardian, "Ephemeral Workload Security in Cloud Environments," [Online]. Available: <https://www.gitguardian.com/nhi-hub/ephemeral-workload-security-in-cloud-environments>
- [2] Ashish Kumar Mishra, et al., "A survey on optimal utilization of preemptible VM instances in cloud computing," ACM Digital Library, 2018. [Online]. Available: <https://dl.acm.org/doi/abs/10.1007/s11227-018-2509-0>
- [3] Prateek Sharma, et al., "Portfolio-driven Resource Management for Transient Cloud Servers," ACM Digital Library, 2017. [Online]. Available: <https://dl.acm.org/doi/10.1145/3084442>
- [4] Eli Cortez, et al., "Resource Central: Understanding and Predicting Workloads for Improved Resource Management in Large Cloud Platforms," ACM Digital Library. [Online]. Available: <https://dl.acm.org/doi/10.1145/3132747.3132772>
- [5] Feng Yan, et al., "Optimizing Power and Performance Trade-offs of MapReduce Job Processing with Heterogeneous Multi-core Processors," IEEE, 2014. [Online]. Available: <https://ieeexplore.ieee.org/document/6973747>
- [6] Shuo Liu, et al., "Profit Aware Load Balancing for Distributed Cloud Data Centers," IEEE, 2013. [Online]. Available: <https://ieeexplore.ieee.org/document/6569848>
- [7] Abhishek Verma, et al., "Large-scale cluster management at Google with Borg," ACM digital library, 2015. [Online]. Available: <https://dl.acm.org/doi/10.1145/2741948.2741964>
- [8] Xiao Zhang, "CPI2: CPU performance isolation for shared compute clusters," The University of Kansas. [Online]. Available: <https://www.ittc.ku.edu/~heechul/courses/eecs750/S14/slides/W4-CPI2-sid.pdf>
- [9] Edo Liberty, et al., "Elastic Machine Learning Algorithms in Amazon SageMaker," ACM Digital Library, 2020. [Online]. Available: <https://dl.acm.org/doi/10.1145/3318464.3386126>
- [10] Haoyu Zhang, et al., "SLAQ: Quality-Driven Scheduling for Distributed Machine Learning," arxiv, 2018. [Online]. Available: <https://arxiv.org/abs/1802.04819>



©2025 by the Authors. This Article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>)