Scholar Publishing Group
International Journal of Big Data Intelligent Technology
https://doi.org/10.38007/IJBDIT.2026.070208
ISSN 2790-0932 Vol. 7, Issue 2: 73-83

Dynamic Resource Allocation and Energy Optimization for Virtualized Data Center

Zhu Xu

Electrical and Computer Engineering, Northwestern University, Evanston 60208, IL, USA

Keywords: Virtualized Data Center; Dynamic Resource Allocation; Virtual Machine Consolidation; Host Power Consumption; CPU Resource Shortage; Trajectory-Driven Simulation

Abstract: To address the issue of virtualized data centers misinterpreting short-term CPU fluctuations as sustained overload during host consolidation, leading to frequent migrations and CPU resource shortages, this study proposes a tiered net-energy-benefit strategy, Tiered Net-Energy-Benefit (Tiered-NEB). This strategy prioritizes temporary MIPS expansion and reclamation within the original host when short-term CPU pressure occurs. When the host experiences sustained low load, the system considers EMA trends, overall migration costs, future startup costs, and state cycle protection to determine whether consolidation should proceed. Based on Bitbrains fastStorage trajectories, this study compares static First-Fit packing, fixed-threshold consolidation, local regression predictive migration, and Tiered-NEB under conditions of 100 homogeneous hosts, 1200 virtual machines, and a continuous 24-hour workload. Results show that Tiered-NEB's host power consumption model has a cumulative energy consumption of 154.21 kWh, a reduction of 8.43% compared to static First-Fit; the CPU resource shortage incidence rate is 1.85%, with 128 completed migrations and 14 host state transition events. This strategy does not achieve the lowest single power consumption, but it can achieve a relatively stable balance between host operating power consumption, CPU resource supply and scheduling stability.

1 Introduction

With the continuous development of cloud computing, online services and enterprise information systems, data centers have become an important infrastructure for the stable operation of digital businesses. Even when servers in data centers are under low load, they still need to keep their processors, memory, network interfaces and basic software environments running, and therefore will continue to consume power. Virtual machine consolidation can centrally deploy workloads to fewer physical hosts and remove temporarily idle hosts from the resource allocation process. This method provides a feasible path to reduce the energy consumption of host operation. Singh and Walia pointed out that host selection, virtual machine selection, placement decisions and migration overhead will jointly affect the consolidation effect [1].

Existing research has improved virtual machine consolidation in terms of load balancing, trend prediction, multi-resource constraints, and intelligent optimization. Yao et al. proposed an energy-efficient load balancing strategy, which reduces resource fragmentation and unnecessary migration by improving resource matching [2]. Vila et al. introduced load prediction into the consolidation process to help the system identify migration opportunities [3]. Sharma et al. incorporated resources such as CPU, memory, and bandwidth into virtual machine allocation and migration decisions [4]. These studies show that host energy consumption optimization cannot rely on a single threshold or focus on only one resource indicator.

However, in actual resource pools, short-term CPU fluctuations and continuous loads are not the same type of state. If the scheduling system migrates immediately based on a single utilization exceeding the limit, short-term load spikes may be treated as continuous overloads. If a host is shut down only because of the current low load, subsequent load recovery may lead to startup and migration again. From the operational scenario focused on in this study, frequent adjustments not only consume network and CPU resources, but also make it more difficult for the resource pool to remain stable. Radi et al. have already considered the number of migrations as an important factor in consolidation decisions [5]. Ma et al. also pointed out that migration timing, migration method, and migration overhead will jointly affect the energy consumption optimization effect of the data center [6]. Therefore, resource scheduling needs to set clear control boundaries between energy saving and operational stability.

Predictive and intelligent methods can identify potential resource pressures in advance, but the system still needs to make judgments based on actual operating costs. Chen et al. incorporated power consumption and thermal factors into virtual machine scheduling to address multi-objective resource optimization problems [7]. Aghasi et al. used distributed reinforcement learning to conduct energy-efficiency-aware virtual machine placement [8]. Chen et al. proposed a combined trend integration strategy and emphasized that the system needs to refer to both current and future resource utilization information [9]. Moocheet et al. used sensor data and learning models to improve energy consumption and thermal control in virtual machine placement [10]. Lin et al. studied energy-efficiency-aware virtual machine placement from the perspective of an overall thermal model [11]. These studies collectively demonstrate that there is a clear trade-off between energy consumption optimization, resource supply, and operating costs.

Based on the above understanding, this study focuses on the dynamic allocation of CPU resources, virtual machine migration and low-load host consolidation within a single virtualized data center, and proposes a tiered net-energy-benefit strategy, Tiered Net-Energy-Benefit, abbreviated as Tiered-NEB. This study does not discuss cross-regional scheduling, GPU heterogeneous computing, container orchestration, cooling systems, PUE or carbon emission trading, but focuses on the host operating power consumption, CPU resource supply and scheduling stability in a fixed resource pool. Tiered-NEB divides resource control into two levels: Tier-1 is for short-term CPU spikes, implementing temporary MIPS expansion and reclamation within the original host; Tier-2 is for continuously low-load hosts, combining trend estimation, whole machine migration cost, future startup cost and host state protection conditions to determine whether to consolidate. The study by Manikandan et al. also shows that virtual machine consolidation needs to consider both energy consumption benefits and additional overhead [12].

To verify the practical performance of this strategy, this study constructs a trajectory-driven simulation scenario using Bitbrains fastStorage trajectories, and fixes the selection of 1200 virtual machines, a continuous 24-hour workload, and 100 homogeneous physical hosts. The experiment sets up four strategies: static First-Fit bin packing, fixed threshold integration, local regression predictive migration, and Tiered-NEB [13]. The system compares five aspects: cumulative energy consumption of the host power consumption model, average number of active hosts, CPU resource shortage incidence, number of completed migrations, and number of host state transition events.

This study does not aim to achieve optimal values for any single strategy across all metrics, but rather examines whether hierarchical control can achieve a more reasonable balance between energy saving, resource supply, and scheduling disturbances [14].

2 Research Design and Methods

2.1 Workload and Simulation Resource Pool

This study selected the fastStorage trajectory from the Bitbrains GWA-T-12 dataset as the workload. To reduce the impact of random sample selection on the results, the system sorted the virtual machine files in ascending numerical order and fixed the first 1200 virtual machine files. The experiment start timestamp was set to 1376314846, and the system continuously captured 86400 seconds of records, forming a 24-hour workload window. A fixed random seed of 42 was used to handle parallel cases during the resource selection process, ensuring consistent decision-making order under the same configuration.

This study constructed a simulation resource pool consisting of 100 homogeneous physical hosts. Each host is configured with 32 processor cores, each with a computing power of 2500 MIPS, 128 GB of memory, and a network bandwidth of 10 Gbps. At the start of the simulation, 75 hosts were active, and 25 hosts were in standby mode. The system used CloudSim's native VmAllocationPolicySimple to complete the initial First-Fit compact binning. The system uses a uniform scheduling cycle of 300 seconds and updates virtual machine requests, host utilization, resource adjustment results, and structured operation records in each cycle.

Let $v$the initial contracted CPU quota of the virtual machine be $a_{v,base}$and the CPU utilization at time t be $r_v(t)$, then its CPU request volume is defined as:

$$d_v(t) = r_v(t)\, a_{v,\text{base}}$$

The system does not include virtual machines that have not yet generated valid records in the current cycle's running set, nor does it include virtual machines whose records have ended in the denominator of the resource shortage rate. Tier-1 can only temporarily adjust the available capacity of virtual machines within the range allowed by the remaining CPU resources of the host machine; therefore, this module will not change the total CPU capacity of the resource pool. The experimental environment and policy configuration are shown in Table 1.

Table 1 Experimental Environment and Strategy Configuration

Parameter categories	Parameter settings
Resource pool size	100 homogeneous physical servers; 32 cores × 2500 MIPS; 128 GB RAM; 10 Gbps
initial state	75 active hosts, 25 deactivated standby hosts
workload	Bitbrains fastStorage; fixed 1200 VMs; continuous 24 hours
Samples and time windows	VM IDs in ascending order; start timestamp 1376314846; random seed 42
Scheduling granularity	300 s
High/Low Load Threshold	0.85 / 0.30
Continuous window	3 cycles, or 15 minutes
EMA parameters	Smoothing coefficient 0.3; recovery rate 5%/h; time domain upper limit 8 h
migration approximation	Net bandwidth 1240 MB/s; combined additional power consumption 10 W
Host state protection	4-hour power status cycle protection window

2.2 Tiered-NEB Strategy

Tier-1 distinguishes between short-term CPU fluctuations and sustained overload. The system maintains a CPU utilization window of 3 scheduling cycles (15 minutes) for each host. If the host CPU utilization exceeds 0.85, but this state has not lasted for 15 minutes, the system will not directly trigger host-level migration. If the CPU utilization of a virtual machine within the host exceeds 0.85, and the host still has available MIPS, the system will provide temporary computing resources for that virtual machine.

To prevent temporary capacity expansion from occupying host capacity for extended periods, the system saves an immutable initial MIPS quota when a virtual machine is created. When the current request volume of a virtual machine falls below 80% of the initial quota, and the virtual machine is in a temporary capacity expansion state, the system restores its initial capacity.

$$d_v(t) \leq 0.8\, a_{v,\text{base}}$$

For sustained high loads exceeding 15 minutes that cannot be absorbed by local expansion, the system enters a unified resource protection process: It prioritizes migrating out the virtual machine with the highest requested MIPS, selecting a target host from the active hosts that still meets CPU and memory constraints after migration and has low power consumption; if no suitable target host exists, a backup host is started and shut down. Only after startup is complete will this host participate in subsequent resource allocation.

Tier-2 only applies to consolidation decisions for hosts with consistently low loads. Let the historical utilization sequence of a host be $\{u_1,u_2,\ldots,u_k\}$, and its exponential moving average be calculated using the following formula:

$$\text{ema}_i = \alpha u_i + (1-\alpha)\text{ema}_{i-1}, \quad \alpha = 0.3$$

When the EMA is above 0.30, the system does not enter the shutdown assessment; when the EMA remains in the low load range, the system estimates the time it takes for the host to reach the 50% utilization safety line based on a recovery rate of 5% per hour, and limits it to within 8 hours.

$$T_{\text{stable}} = \min\left(\frac{0.50 - \text{ema}}{0.05}, 8\right)$$

For a candidate host h to be shut down, the system considers all virtual machines on that host as a complete decision object for migration. Let be the set of virtual machines on the source host, $V_{\text{h}}M_v$be the configured memory of virtual machine v, and the net migration bandwidth be 1240 MB/s. Then, the migration time and the additional energy consumption for the entire machine migration are approximately:

$$T_{\text{mig},v} = \frac{M_v}{1240}, \quad \hat{E}_{\text{mig},h} = \sum_{v\in V_h} \frac{T_{\text{mig},v}}{3600} \times 10$$

The host startup time was set to 90 seconds, and the startup power consumption was calculated at 222 W, corresponding to a startup cost of 5.55 Wh. Using 58.4 W as the low-load decision parameter for this experiment, the net benefit proxy for host shutdown is:

$$\widehat{\text{NEB}}_h = T_{\text{stable}} \times 58.4 - \hat{E}_{\text{mig},h} - 5.55$$

The system only switches the host to a shutdown standby state when the net benefit is positive, the host meets the 4-hour state cycle protection condition, and all virtual machines have been successfully migrated out. This net benefit is only used to filter out uneconomical shutdown actions and is not directly added to the cumulative energy consumption of the host power consumption model in the results table. The Tiered-NEB hierarchical resource control process is shown in Figure 1.

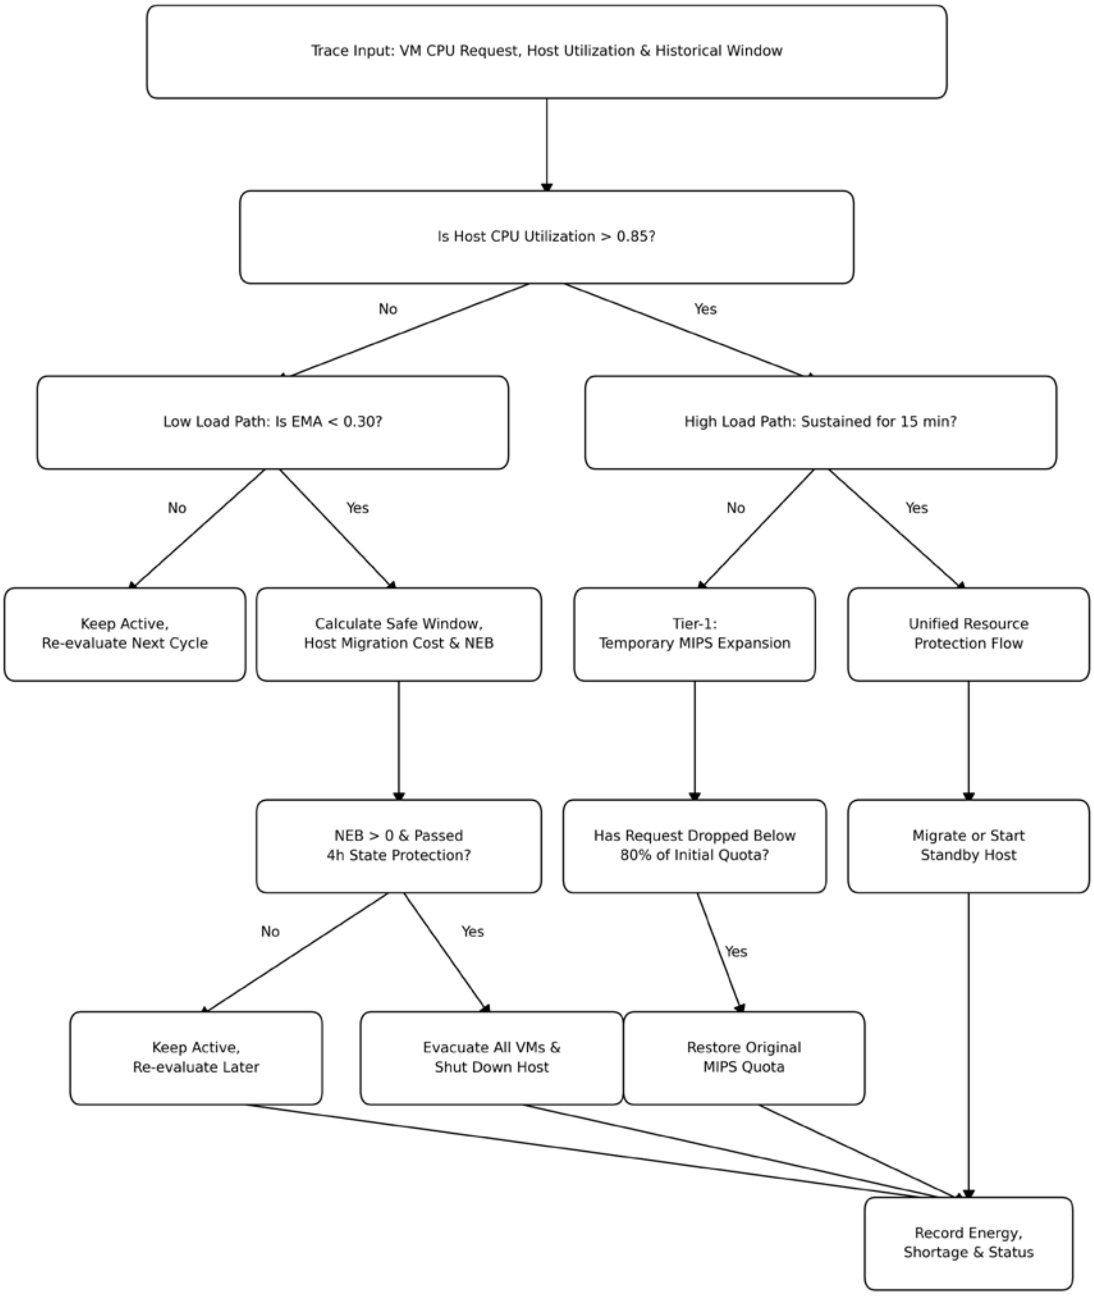

Figure 1 Tiered-NEB hierarchical resource control process

2.3 Comparison Strategy and Evaluation Indicators

B1 employs a static First-Fit packing strategy. After the experiment begins, B1 maintains 75 active hosts and 25 hosts in a standby state, without performing virtual machine migrations, host state adjustments, or temporary MIPS expansion. B2 employs a fixed-threshold consolidation strategy. When a host's CPU utilization exceeds 0.85 or falls below 0.30 for a single instance, B2 immediately triggers a migration or consolidation action. B2 does not use persistent windows, net benefit judgments, or state loop protection. B3 employs a local regression predictive migration strategy. B3 uses local regression and median absolute difference (MAD) to identify potential overloads, with a safety margin of 1.2. B3 does not set Tier-1 local expansion and does not use overall system migration costs, startup costs, or state loops as explicit constraints.

The four strategies use the same workload, initial packing, host configuration, scheduling cycle, and power consumption model. The system does not allow a single virtual machine to re-enter the migration queue before a previous migration is complete. The destination host must simultaneously meet CPU and memory capacity constraints. Standby hosts in the startup phase are not considered available targets until startup is complete. The system only increases the migration count after the migration is complete and the source host resources have been released.

This study defines the cumulative output of the host power consumption model as the cumulative host energy consumption:

$$E_{\text{host}} = \frac{\sum_{h\in H} E_h^{\text{model}}}{3.6 \times 10^6}$$

This metric is used to compare host-side operating power consumption under different strategies, and is not interpreted as including all additional costs of cooling, power distribution, migration, and startup in the total data center energy consumption. CPU resource shortage incidence rate is defined as:

$$\text{CSR} = \frac{\sum_{t\in T} \sum_{v\in V_t} I\left(d_v(t) > a_v(t)\right)}{\sum_{t\in T} |V_t|} \times 100\%$$

This represents the set of virtual machines that are $V_t$always *t*running. This metric only indicates the proportion of VMs with insufficient CPU supply and is not equivalent to the business SLA default rate. The number of host state transition events represents the total number of completed startup and shutdown events.

2.4 Experimental Records and Control Methods

To ensure the comparability of the four strategies, the system records the time, number of active hosts, cumulative migration count, number of host state transition events, average CPU utilization, CPU resource shortage occurrence rate, and cumulative energy consumption of the host power consumption model after each 300-second scheduling cycle. B1, B2, B3, and Tiered-NEB use the same resource pool, trajectory window, initial binning, and power consumption model. The different methods only change the logic for short-term high-load processing, low-load consolidation, and host shutdown triggering.

The system only counts completed virtual machine migration actions, excluding migration candidates or incomplete requests. The system accumulates host state transition events based on completed startup and shutdown actions. B2 and B3 currently retain the total number of state transition events. Tiered-NEB's initial active host count was 75, the final active host count was 65, and the number of state events was 14; these three records satisfy the state conservation relationship. This process avoids conflating planned actions, completed actions, and changes in the number of hosts into a single metric.

The system aggregates the cumulative outputs of each physical host power consumption model into a cumulative energy consumption model for host power consumption, used to compare differences in host operating power consumption caused by resource scheduling. Tier-2 uses migration-related additional power consumption and startup costs for screening the net benefit of shutting down the entire machine, but the system does not include these two costs as separate ledger items in Table 2 to avoid duplicate measurement under the current monitoring scope. The system uses the VM-cycle of running virtual machines as the denominator for the CPU resource shortage occurrence rate, ensuring that different strategies use a consistent statistical caliber.

3. Experimental Results and Analysis

3.1 Comparison of Overall Results

Table 2 presents the overall results for the four strategies under the conditions of a fixed 1200 virtual machines, a continuous 24-hour workload, and 100 physical hosts. To check for any significant discrepancies between energy consumption and host size, the "equivalent average power consumption per active host" is added to the table as a consistency verification item. This value is obtained by dividing the cumulative host energy consumption by the average number of active hosts and the 24-hour duration.

Table 2 Overall experimental results of the four strategies

Method	Cumulative host power consumptio n/kWh	Average number of active hosts/units	Single unit equivalent average power consumptio n/W	CPU resource shortage incidence rate / %	Number of migrations completed	Number of state transition events
B1 Static First-Fit	168.41	75.0	93.56	2.15	0	0
B2 Static Threshold	142.54	52.3	113.56	8.70	2845	312
B3 Local Regression	146.12	54.8	111.10	6.45	2130	275
Proposed Tiered-NEB	154.21	64.5	99.62	1.85	128	14

B1 maintained 75 active hosts throughout, resulting in the highest cumulative host energy consumption of 168.41 kWh. B2 and B3 reduced the average number of active hosts through stronger host consolidation, lowering their cumulative host energy consumption to 142.54 kWh and 146.12 kWh, respectively. Tiered-NEB's cumulative host energy consumption was 154.21 kWh, a decrease of 14.20 kWh from B1, representing an 8.43% reduction. Tiered-NEB's average number of active hosts was 64.5, a decrease of 10.5 from B1, indicating that this strategy freed up some hosts with consistently low load.

In terms of host compression, B2 has the lowest average number of active hosts at 52.3; B3 has 54.8; and Tiered-NEB has 64.5. B2 and B3 run fewer hosts, but each active host bears a higher average load, resulting in equivalent average power consumption of 113.56 W and 111.10 W respectively. B1 runs more hosts, with lower average power consumption per host, but its large number of hosts results in the highest cumulative energy consumption. Tiered-NEB's equivalent average power consumption per host is 99.62 W, falling between static deployment and aggressive consolidation, a result consistent with the resource pool's operating state.

Figure 2 illustrates the trade-offs among the four strategies using cumulative host energy consumption and CPU resource shortage incidence rate. B2 and B3 fall within the lower energy consumption, higher shortage range; B1 falls within the higher energy consumption, lower shortage range; Tiered-NEB is in the middle energy consumption range, but has the lowest shortage incidence rate. This result does not show a scenario where all metrics are optimal simultaneously, which aligns well with the multi-objective constraints in virtualization resource scheduling. It is important to note that the lower cumulative energy consumption of B2 and B3 does not alone indicate better overall scheduling performance, as frequent migrations and state adjustments also consume network, CPU, and operational resources.

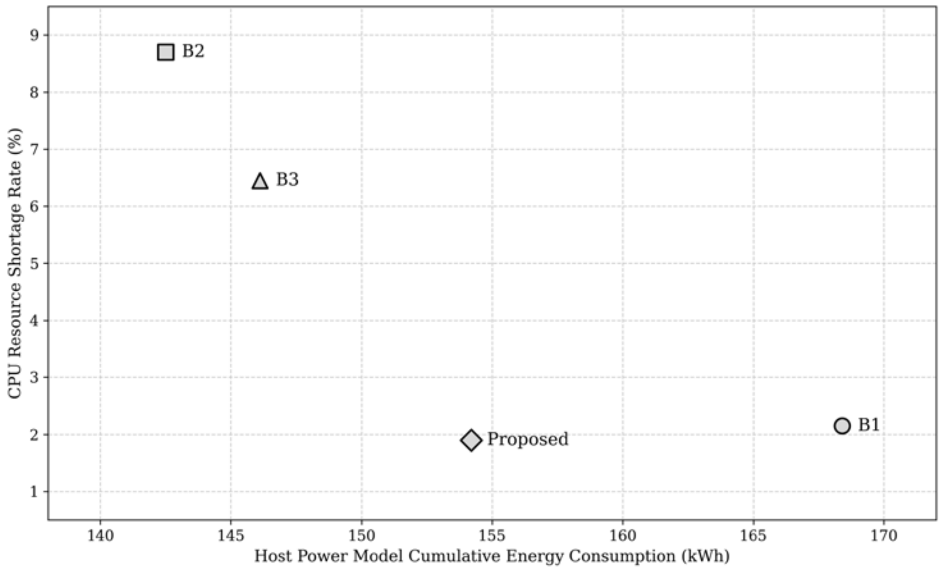

Figure 2. Trade-offs between host power consumption and CPU resource shortage for four strategies.

3.2 CPU Resource Shortage and Scheduling Disturbance Analysis

B2 had the highest CPU resource shortage rate at 8.70%. B3's shortage rate was 6.45%, lower than B2 but still significantly higher than B1 and Tiered-NEB. B2 and B3 reduced host power consumption by rapidly compressing the size of active hosts, but the smaller resource margin also made the system more susceptible to situations where CPU requests could not be met in a timely manner when the load changed.

The CPU resource shortage rate of Tiered-NEB was 1.85%, a decrease of 0.30 percentage points compared to B1, 6.85 percentage points compared to B2, and 4.60 percentage points compared to B3. While static B1 runs a larger number of hosts, it lacks the ability for local expansion and dynamic migration, thus the system may still experience resource fragmentation and localized supply shortages. Tiered-NEB retains local adjustment space based on limited host consolidation, which is a key reason why its shortage rate is lower than that of B1.

The number of migrations reflects the scheduling disturbances caused by different strategies. B2 completed 2845 migrations in 24 hours, averaging 118.54 migrations per hour; B3 completed 2130 migrations, averaging 88.75 migrations per hour; and Tiered-NEB completed only 128 migrations, averaging 5.33 migrations per hour. Compared to B2, Tiered-NEB reduced migrations by 2717, a decrease of 95.50%; compared to B3, Tiered-NEB reduced migrations by 2002, a decrease of 93.99%. From the actual operation of the resource pool, this difference means that Tiered-NEB is less likely to convert short-term fluctuations into cross-host adjustments.

The host state transition results show a similar trend to the number of migrations. B2 and B3 recorded 312 and 275 state transition events, respectively, while Tiered-NEB only recorded 14. Tiered-NEB initially had 75 active hosts, which decreased to 65 at the end of the simulation. If the 14 counts represent the total number of completed startup and shutdown events, then the system corresponds to 12 host shutdowns and 2 host startups, and the number of hosts satisfies the conservation relationship. B2 and B3 currently only retain the total number of state events, therefore this study will not further separate the startup and shutdown counts for them.

3.3 Typical Operational Events

In a short-term high-load scenario, the CPU utilization of Host-14 rose to 88.2% around 10800 s of simulation time. This high-load state did not last for the 15-minute window, so Tier-1 did not trigger a host-level migration, but instead implemented temporary MIPS expansion for VM-204. Subsequently, the CPU utilization of Host-14 dropped back to 72.4%. This process demonstrates that the system can prioritize handling short-term CPU pressure within the original host, rather than directly translating each utilization peak into a cross-host migration.

In a sustained low-load scenario, Host-62 remained under low load for 15 minutes around simulation time 67500 s. The system estimates the low-load safety window for this host to be 4.8 hours and calculates a projected host operating saving of 280.3 Wh based on a low-load decision parameter of 58.4 W. This host contains 4 virtual machines, with a total migration cost of 0.16 Wh and a future startup cost of 5.55 Wh. Therefore, the system obtains:

$$\widehat{\mathrm{NEB}} = 280.3 - 0.16 - 5.55 = 274.59\,\mathrm{Wh}$$

Host-62 showed positive net gains and passed the 4-hour state cycle protection check. The system subsequently completed the migration of four virtual machines and put Host-62 into a shutdown standby state. This event demonstrates that Tier-2 does not shut down a host solely based on a single low load, but rather considers trends, migration costs, startup costs, and protection conditions simultaneously.

4. Discussion

Experimental results show that the four strategies have different resource management orientations. B1 maintains a static resource layout, avoiding migration and state adjustments, but B1 cannot release underutilized hosts, thus resulting in the highest cumulative host energy consumption. B2 rapidly compresses host size through a single threshold, achieving the lowest cumulative host energy consumption, but B2 also results in the highest CPU resource shortage rate and the largest migration volume. B3 uses local regression and MAD to identify potential overload, and its shortage rate and migration volume are lower than B2, but B3 still maintains a high intensity of resource reconfiguration.

The key feature of Tiered-NEB is not achieving the lowest single power consumption metric, but rather controlling resource shortages and scheduling disturbances while reducing host operating power. Tier-1 prioritizes limiting short-term CPU load to the original host, preventing single utilization peaks from directly entering the migration path. Tier-2 checks the sustainability of low loads, overall migration costs, future startup costs, and state cycle protection conditions before host shutdown, thus reducing the likelihood of brief low loads being amplified into frequent host state adjustments.

From an engineering application perspective, Tiered-NEB is suitable for virtualization resource pools that need to balance host power consumption, CPU resource allocation, and operational stability. This strategy does not rely on complex black-box models; the system can interpret migration and shutdown decisions through utilization windows, MIPS requests, EMA values, migration costs, startup costs, and state protection records. For operations and maintenance personnel, this traceability helps identify abnormal migrations, examine shutdown decisions, and adjust threshold settings.

The conclusions of this study have a clearly defined scope. The cumulative host energy consumption in Table 2 comes from the host power consumption model and is mainly used to compare the host-side operating power consumption under different scheduling strategies. The system has included migration-related power consumption and startup costs in Tier-2 shutdown screening, but these two costs have not been included as independent ledger items in the total energy consumption comparison of the four strategies. Therefore, this study does not interpret Table 2 as the total data center energy consumption including all costs of cooling, power distribution, migration, and startup. This study only uses a single fastStorage trajectory, a fixed 1200 virtual machines, and a continuous 24-hour window. Future studies can use more load windows to test parameter stability while keeping the main strategy unchanged.

5. Conclusion

This study proposes a Tiered-NEB dynamic resource allocation strategy to coordinate host operating power consumption, CPU resource supply, and scheduling stability in virtualized data centers. Tier-1 reduces unnecessary migrations by identifying short-term CPU spikes and implementing temporary MIPS expansion. Tier-2 determines whether a host should be taken offline based on low load trends, overall migration costs, future startup costs, and state cycle protection.

Under conditions of a fixed 1200 virtual machines, a continuous 24-hour trajectory, and 100 homogeneous hosts, Tiered-NEB achieved a cumulative host energy consumption of 154.21 kWh, a reduction of 8.43% compared to the static First-Fit strategy. This strategy reduced the average number of active hosts from 75.0 to 64.5 and controlled the CPU resource shortage rate to 1.85%. The system completed 128 migrations and recorded 14 host state transition events. Compared to fixed threshold and local regression predictive migration, Tiered-NEB did not achieve the lowest cumulative host energy consumption, but it significantly reduced CPU resource shortages, migration counts, and host state adjustment frequency.

The research results indicate that virtualized data center resource scheduling should not solely aim to reduce the number of active hosts. By incorporating short-term CPU fluctuations, sustained low loads, overall machine migration costs, and host state protection into the same control process, a more controllable balance can be achieved between host operating power consumption, CPU resource availability, and scheduling stability.

References

[1] Singh J, Walia N K. A comprehensive review of cloud computing virtual machine consolidation[J]. IEEE Access, 2023, 11: 106190-106209.

[2] Yao W, Wang Z, Hou Y, Zhu X, Li Vila S, Guirado F, Lérida J L. Cloud computing virtual machine consolidation based on stock trading forecast techniques[J]. Future Generation Computer Systems, 2023, 145: 321-336.

[3] Ma, W. (2026). Reliable Data Infrastructure Supported by Automated Data Pipelines. Engineering Advances, 6(2).

[4] Li, J. (2026). Research on the Path and Efficiency of Empowering Accurate Advertising Delivery with Data Infrastructure.

[5] Zhang, Y. (2026). Research on LLM-Driven Intelligent Architecture Design and Autonomy Mechanism Construction for Cloud-Native Control Planes.

[6] Guo, X. (2025, April). Research on Financial Trading Algorithms Based on Deep Reinforcement Learning. In 2025 IEEE 3rd International Conference on Control, Electronics and Computer Technology (ICCECT) (pp. 1711-1715). IEEE.

[7] Liu, X. (2026). Construction of Consumer Data Privacy Protection System Based On Blockchain Technology. Procedia Computer Science, 282, 2082-2091.

[8] Guo, X. (2025, March). Research on Blockchain-Based Financial AI Algorithm Integration Methods and Systems. In 2025 IEEE International Conference on Electronics, Energy Systems and Power Engineering (EESPE) (pp. 816-821). IEEE.

[9] *Li, J. (2026). Analysis of Advertising Creativity Generation and User Response Mode Based on AIGC.*
[10] *Zhang, Y. (2026). A Framework for LLM-Based Semantic Configuration Understanding and Automated Change Generation in Microservice Orchestration.*
[11] *Liu, X. (2026). Research on the Application of Artificial Intelligence in Consumer Privacy Data Protection. Procedia Computer Science, 279, 956-965.*
[12] *Zhang, Y. (2026). Exploration of Enterprise Big Data Microservice Architecture Based on Domain-Driven Design (DDD). Procedia Computer Science, 282, 1994-2003.*
[13] *Wang, Z. (2026). Research on Data Analysis and Quality Prediction of Manufacturing Processes Based on Improved Deep Learning Algorithms. Procedia Computer Science, 281, 1328-1337.*
[14] *Wang, N. (2026). Construction and Application of Interpretable Enterprise Performance Prediction Model Based on Multi-Source Operational Data.*