Multi-Cloud Management Architecture Design and Disaster Recovery Strategy for High Availability

Wenchao Li, Guang Ma, Wenchong Fang*, Xiqi He and Jin Li

China Southern Power Grid Corporation Limited Power Dispatch Control Center, Guangzhou 510663, Guangdong, China
E-mail: fangwc@csg.cn
**Corresponding Author*

Received 11 September 2025; Accepted 21 October 2025

Abstract

With the growing adoption of multi-cloud computing for mission-critical applications, ensuring high availability and designing resilient disaster recovery strategies have become paramount. Existing solutions often lack a unified architecture for managing heterogeneous resources and rely on reactive, rule-based recovery mechanisms. To address these gaps, this paper proposes the High-Availability Multi-cloud Management Model with Intelligent Orchestration and Disaster Recovery (HMM-IODR). The model integrates cross-cloud resource orchestration, real-time fault detection, automated failover, and consistent data replication mechanisms.A virtual multi-cloud testbed was implemented to emulate regional failures across provider-specific scenarios, with training and evaluation data derived from public M-Lab NDT measurements. Several disaster recovery strategies were evaluated, including multi-live deployment, hot standby, cold standby, and intelligent CDN-based traffic redirection. Experimental results showed that the multi-live deployment strategy yielded a Recovery Time Objective (RTO) of 3.45 seconds and maintained latency below 188.7 ms. This RTO represents an

Journal of Cyber Security and Mobility, Vol. 14_5, 1173–1198.
doi: 10.13052/jcsm2245-1439.1456

88.5% reduction compared to the Cold Standby strategy under identical test conditions, demonstrating effective scalability. Further analysis of scalability metrics confirmed the effectiveness of the HADRM model in mitigating fault impact and enhancing overall system availability. This study provides a practical framework for building scalable, secure, and fault-tolerant multi-cloud information systems, contributing to the advancement of cloud-native distributed architectures.

Keywords: Scalable distributed information systems, multi-cloud architecture, high availability, disaster recovery, content delivery networks (CDN), scalability metrics, cloud-native architectures, information security.

1 Introduction

The accelerated promotion of enterprise digital transformation has driven the evolution of cloud computing from single-cloud services to multi-cloud collaboration. Multi-cloud architectures effectively mitigate single-point dependencies and service interruption risks by deploying applications and data across multiple cloud service providers, and they are gradually becoming the fundamental infrastructure supporting critical business systems [1]. In domains such as finance, transportation, and energy, where business continuity is paramount, high availability and disaster recovery capabilities have become the core objectives of multi-cloud management [2].

However, the heterogeneity, distribution, and dynamic network conditions of multi-cloud environments introduce unprecedented complexity to resource scheduling, fault detection, data consistency, and disaster recovery strategies [3]. Traditional research has primarily focused on high availability within a single cloud or across regions of the same vendor, but lacks unified scheduling mechanisms for cross-cloud heterogeneous resources. As a result, resource utilization remains low and deployment strategies inflexible. Similarly, disaster recovery solutions are often reactive, relying on manual or rule-based mechanisms, which lack intelligent analysis and real-time decision-making. Moreover, existing data consistency solutions frequently depend on centralized storage or third-party trust models, and thus cannot provide decentralized verification or trusted synchronization across cloud nodes [4–5].

Multi-cloud database systems also face challenges such as data fragmentation and latency, with traditional methods struggling to ensure interoperability. To address these issues, Li et al. proposed a federated cloud/edge (FCE)

architecture capable of coordinating distributed multi-task processing across multiple cloud sites. By incorporating a voting mechanism and model integration strategy, the FCE architecture improved communication efficiency and training accuracy while maintaining data locality, making it particularly suitable for privacy-sensitive applications such as medical image analysis [6]. Zhang et al. developed a multi-objective dynamic resource allocation model and improved NSGA-II algorithm to address the high dynamism of Industry 4.0 manufacturing systems. By enhancing population diversity and global search, their method significantly improved allocation efficiency and reliability [7]. Similarly, Addya et al. highlighted the importance of validation in production scheduling optimization and proposed an intelligent verification system based on discrete event simulation (DES), confirming the correctness of deep reinforcement learning–based scheduling solutions [8].

Globa L. et al. addressed the challenge of resource allocation and provider selection in multi-cloud environments by employing ontology-based methods and integrating reinforcement learning with multi-objective evolutionary algorithms. Their model standardized domain concepts and enabled effective provider selection, improving automation and decision-making efficiency in high availability and disaster recovery [9]. Zhang et al. further proposed a multi-cloud edge collaborative multi-cluster system architecture based on container technology, incorporating dynamic resource allocation and multi-path transmission mechanisms. Simulation results confirmed its feasibility and improved service responsiveness in integrated air–space–ground networks [10]. In the context of 6G, Kim B. et al. emphasized the importance of distributing massive computing demands across cloud–edge infrastructures to meet high throughput and low latency requirements, highlighting the value of multi-cluster orchestration for resource allocation and collaborative scheduling [11]. Benmerar T. Z. et al. introduced an AI-integrated multi-domain edge orchestration architecture with visualization tools and modularized service deployment, validated in virtual tourism scenarios, demonstrating improved immersive service execution [12]. Taghinezhad-Niar et al. applied digital twin technology to intelligent manufacturing, proposing a data-driven model synchronized with real systems via Simio software. Their experiments confirmed accuracy within 3% for key performance indicators, underscoring the role of high-quality data in real-time decision-making [13].

In summary, extensive research has been conducted on high availability, disaster recovery optimization, resource scheduling, and multi-cluster orchestration in multi-cloud environments. Existing studies have contributed effective solutions leveraging AI, containerization, security management,

and distributed data consistency. However, a unified architecture design for heterogeneous multi-cloud environments remains lacking, hindering cross-cloud collaboration and intelligent orchestration of resources.

To address these gaps, this study proposes a High-Availability Multi-cloud Management Model with Intelligent Orchestration and Disaster Recovery (HMM-IODR). The architecture integrates a blockchain-based data consistency mechanism, a primary–backup node capability evaluation model, an intelligent scheduling algorithm, and the Practical Byzantine Fault Tolerance (PBFT) consensus protocol to provide a systematic solution for ensuring business continuity. Furthermore, a dynamic strategy module supporting multiple recovery mechanisms is designed to adapt to diverse vendor combinations and fluctuating network conditions. This study aims to provide both theoretical foundations and technical support for future exploration and experimental validation in multi-cloud management and disaster recovery.2. Multi-cloud management architecture design for high availability

2 Design of Multi-cloud Management Architecture

To provide a rigorous foundation for our work, this section formally defines the High-Availability Multi-cloud Management Model with Intelligent Orchestration and Disaster Recovery (HMM-IODR). HMM-IODR is conceptualized as a holistic framework that systematically addresses the challenges of multi-cloud resilience. Its definition spans three distinct levels: a guiding conceptual model, a concrete system architecture, and a set of specified interface protocols.

At the highest level of abstraction, the HMM-IODR conceptual model establishes the core principles for resilient multi-cloud management. The model is founded on a layered abstraction that decouples the system into Infrastructure, Orchestration, and Intelligence layers, thereby simplifying the management of heterogeneous cloud resources. Building upon this layered foundation, the model champions an intelligence-driven orchestration approach, which replaces static, rule-based mechanisms with dynamic, data-driven algorithms for predictive resource scheduling and optimal disaster recovery. Furthermore, to ensure reliability across disparate environments, the model mandates a decentralized trust and consistency mechanism, which is critical for validating state transitions and guaranteeing data integrity during failover events without a single point of failure.

Translating these principles into a practical blueprint, the HMM-IODR system architecture specifies the tangible software components and their

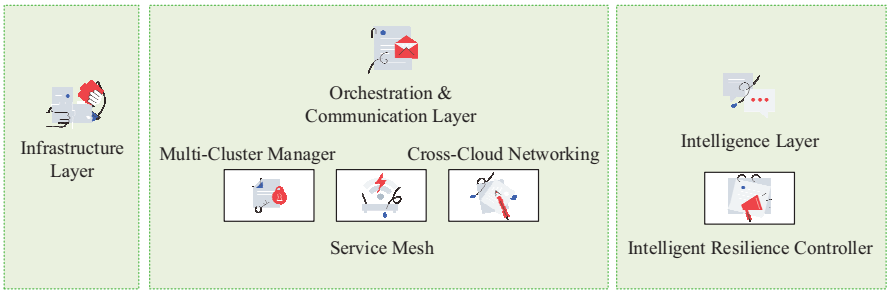

Figure 1 System architecture diagram.

interconnections. While Figure 1 offers a high-level view, a more formal UML component diagram would detail this structure. The architecture's "brain" is the Intelligent Resilience Controller in the Intelligence Layer, which exposes interfaces like IResiliencePolicy and INodeEvaluator to make strategic decisions. This controller directs the Multi-Cluster Manager within the Orchestration Layer, which acts as the system's execution arm. The manager implements the IResourceOrchestration interface to handle application deployment and failover, in turn utilizing the Cross-Cloud Networking and Service Mesh components to manage traffic routing via an IServiceConnectivity interface. Crucially, a dedicated Blockchain Consistency Module provides an IDataConsistency interface, allowing the manager to verify data integrity before finalizing a failover, thus ensuring a trusted recovery process.

To govern these architectural interactions, the framework specifies a set of interface protocols for reliable and standardized communication. These interactions, as detailed in Table 1, define the precise rules of engagement between components. For instance, the Intelligent Resilience Controller issues failover commands to the Multi-Cluster Manager via gRPC or REST calls, while the manager uses gRPC to interact with the Blockchain Consistency Module for data verification. This protocol-driven approach ensures that high-level policies are executed consistently and reliably at the operational level.

Finally, it is essential to delineate the core innovations of the work. The novelty of the HMM-IODR framework lies not in the invention of its underlying technologies, such as blockchain or PBFT, but in their synergistic integration and novel application. The primary contribution is the design of a unified architectural blueprint, HMM-IODR, that cohesively integrates cross-cloud orchestration with intelligent scheduling and decentralized data verification. A second key innovation is the development of an intelligent decision-making engine, based on the service capability evaluation model and

backup selection algorithm presented in this study (Equations 3-8), which transforms disaster recovery from a reactive process into a proactive and optimized one. Ultimately, the core contribution is the integrated framework itself, which is experimentally validated to achieve significant improvements in key disaster recovery metrics like RTO, latency, and success rate within complex multi-cloud scenarios.

Multi-cloud strategy is the mainstream of enterprise digital transformation, and building a management architecture that ensures business continuity and high availability requires systematic and intelligent design [14]. This architecture needs to address management complexity, security challenges, network latency, and data consistency issues in heterogeneous environments. Its core goal is to integrate the resources of different cloud service providers to form a highly resilient, scalable, and automatically responsive whole. Therefore, this study designs a layered, high availability-oriented multi-cloud management architecture, including infrastructure layer, abstraction and orchestration layer, and intelligence and policy layer, as shown in Figure 1.

As shown in Figure 1, heterogeneous resources are managed by the system's management and orchestration modules, which coordinate cross-cloud deployment, optimization and failover while ensuring data consistency. When a user request arrives, the intelligent layer evaluates the cloud node's service capabilities and selects primary and backup nodes that meet QoS requirements. The scheduling algorithm formulates deployment strategies based on real-time state and optimization objectives, and the orchestration layer executes cross cloud deployment and traffic configuration accordingly. At runtime, the data consistency module utilizes blockchain technology to ensure complete synchronization of cloud data replicas. If the monitoring system detects a main node failure or performance drop, the architecture will automatically trigger disaster recovery, and the orchestration layer will smoothly switch traffic to the backup node. After verifying the data through blockchain, the service will be quickly restored to ensure high availability. Figure 2 shows the high-availability multi-cloud management network model.

In Figure 2, the architecture is divided into network user layer, cloud resource layer, and scheduling layer. The network user layer covers diverse end users and devices, and is a consumer of cloud services. The cloud resource layer is composed of heterogeneous resources from different providers, forming a unified resource pool. The multi-cloud orchestration in the scheduling layer is responsible for receiving requests and intelligently

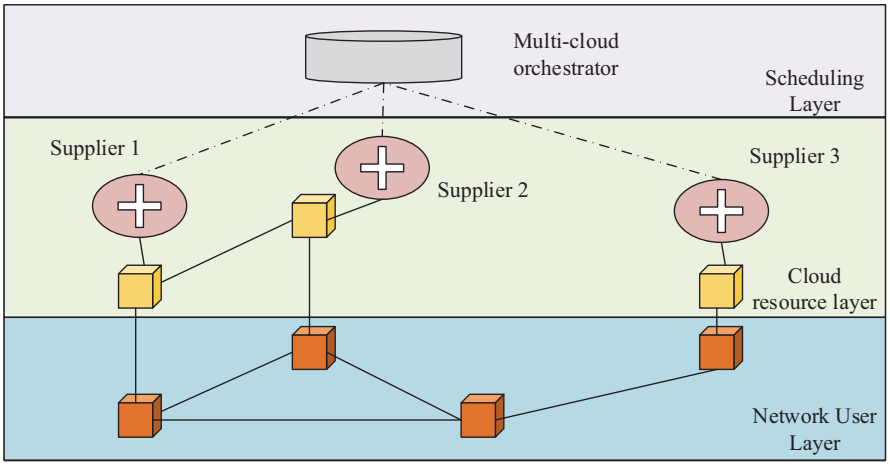

Figure 2 A Multi-cloud management network model for high-availability.

scheduling lower level resources to achieve efficient service deployment and management [15]. However, application orchestration alone does not ensure high availability. In multi-cloud deployments, high availability ultimately depends on preserving data integrity and enforcing cross-cloud consistency during replication and failover. Especially in disaster recovery, it is necessary to ensure consistency between the backup site data and the primary site at a certain point in time before the failure. This study proposes a data consistency assurance module based on blockchain technology. When a data file F needs to be stored and synchronized in a multi-cloud environment, the system first divides it into n standard sized data blocks, namely $F = \{m_1, m_2, \ldots, m_i, \ldots, m_n\}$. Subsequently, to ensure data privacy and security, efficient algorithms such as elliptic curve encryption are used to encrypt each data block m_i, resulting in an encrypted data block λ_i. Then, collision resistant hash functions (such as SHA-256) are used to operate on each encrypted block, generating a fixed length data label γ_i, as shown in equation (1).

$$\gamma_i = H(\lambda_i) \tag{1}$$

In equation (1), H denotes a collision-resistant cryptographic hash function, such as SHA-256, and γ_i represents the resulting data label, which is a fixed-length and unique hash value generated for the i-th encrypted block, λ_i is the i-th encrypted data block, which was previously generated from the original data block. To aggregate the verification information of all data

blocks into a lightweight and unique proof, this module utilizes the Merkle Hash Tree (MHT) structure to hash all data labels γ_i layer by layer, ultimately generating a unique Merkle root R. This root value F can be regarded as an "integrity snapshot" of the entire data file R in a specific state. The generation process can be abstractly represented as shown in equation (2).

$$R = \mathrm{MHT}(\{\gamma_1, \gamma_2, \ldots, \gamma_n\}) \tag{2}$$

In equation (2), R, as a lightweight verification proof, will be recorded in a new block of the blockchain. The top layer of the architecture is the intelligence and strategy layer, which is the "brain" of the entire system, responsible for formulating optimal service deployment and disaster recovery decisions. A cloud node with the same nominal configuration may have significant differences in its actual service capabilities due to underlying hardware aging, resource contention, or network fluctuations. The intelligent layer of the architecture introduces a service capability evaluation module to comprehensively evaluate multidimensional service quality indicators collected from historical services, such as availability, service latency, packet loss rate, etc. [16]. For a set of service quality measurement indicators $\Theta = \{\theta_1, \theta_2, \ldots, \theta_m\}$, the module first determines the importance of each attribute θ_j by calculating its entropy increase $E(\theta_j)$ to the classification result, and then assigns a weight w_j to it, as shown in equation (3).

$$w_j = \frac{E(\theta_j)}{\sum_{j'=1}^{m} E(\theta'_j)} \tag{3}$$

In equation (3), $E(\theta_j)$ reflects the information content of attribute θ_j in distinguishing cloud node service capabilities. After obtaining the weights, for any cloud node cl_m, its comprehensive service capability evaluation value Φ_m can be obtained by weighted sum of its standardized service quality indicators, as shown in equation (4).

$$\Phi_m = \sum_{j=1}^{m} w_j \cdot nor(x_{m,j}) \tag{4}$$

In equation (4), $x_{m,j}$ is the historical performance data of cloud node cl_m on attribute θ_j. $nor(\cdot)$ is the normalization function. Based on this evaluation value, the dynamic scheduling and deployment algorithm of the intelligent layer (such as the algorithm based on online learning or primal duality technology mentioned in the literature) can formulate the global optimal

strategy to maximize platform throughput or service deployment profit while meeting user needs, cloud node resource capacity, and service capabilities constraints.

3 Design and Implementation of Disaster Recovery Strategy

This study has established a unified multi cloud management architecture, laying the foundation for high availability, but a complete system still requires disaster recovery strategies to ensure business continuity. The traditional disaster recovery model is difficult to meet modern business requirements in multi cloud environments. Therefore, further design of dynamic intelligent disaster recovery strategies aims to automatically and reliably switch services to the optimal backup node at the lowest cost and fastest speed, ensuring the integrity of business and data [17]. The multi-level disaster recovery architecture is shown in Figure 3.

In Figure 3, the architecture is monitored and coordinated by the top-level disaster recovery management module. The core production environment adopts the same city dual active disaster recovery mode, deployed in the availability zones (Active AZ1, Active AZ2) of geographically close and physically isolated data centers (DC1, DC2). It ensures strong data

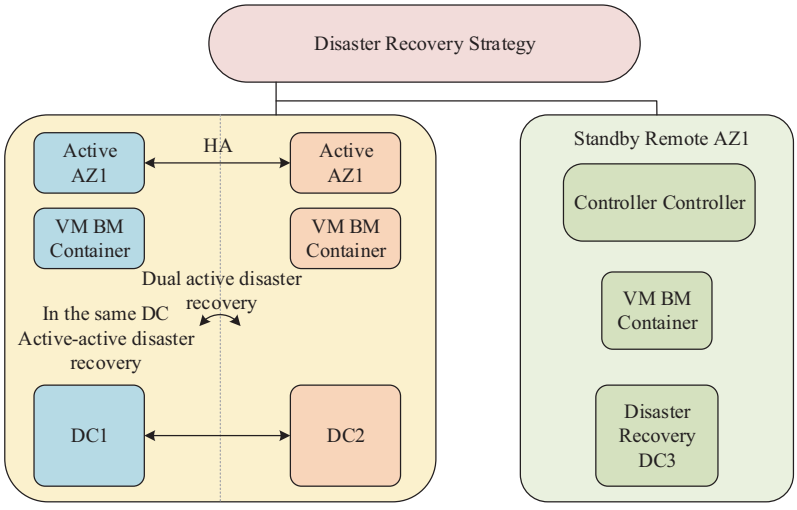

Figure 3 Multi-level disaster recovery architecture.

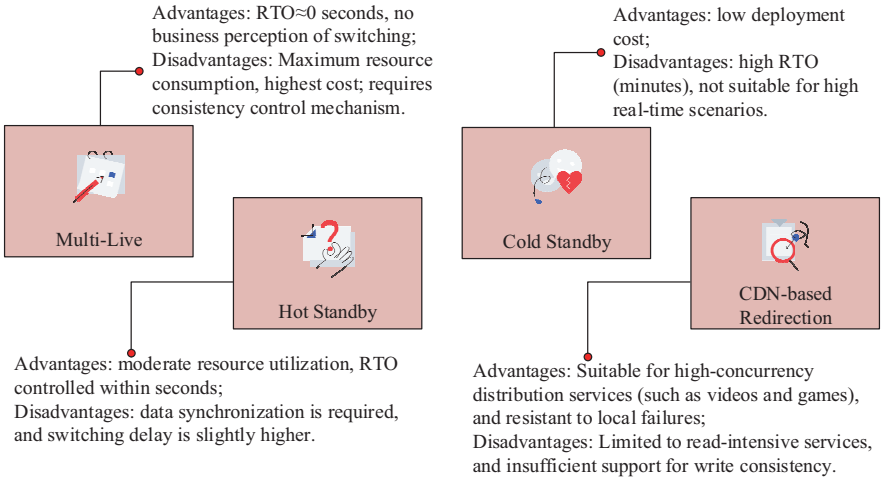

Figure 4 Four disaster recovery modes.

consistency and fast fault switching through the high availability interconnection of the controller and synchronous replication of the underlying storage. To cope with regional disasters, the remote primary backup disaster recovery mode has been introduced, deploying a disaster recovery center (Disaster Recovery DC3) in the remote availability zone (Standby Remote AZ1). By asynchronously replicating and transmitting data, the disaster recovery center can be activated to take over the service in extreme cases. According to different requirements for business continuity and recovery time objectives, four disaster recovery modes are defined, as shown in Figure 4.

To achieve this dynamic classification, the system uses exponential smoothing to continuously evaluate the access frequency of each data file. The estimated value $St_k(t_n)$ of the access popularity of a specific data file k in time slice t_n can be calculated using equation (5).

$$St_k(t_n) = \alpha \cdot y_{t_n} + (1-\alpha) \cdots t_k(t_{n-1}) \tag{5}$$

In equation (5), y_{t_n} is the actual number of visits observed within the current time slot. $St_k(t_{n-1})$ is the estimated heat value of the previous time slot. $\alpha(0 < \alpha < 1)$ is a smoothing coefficient (or attenuation factor) used to control the weight of new and old observations in estimation. A higher α value means that the system is more sensitive to recent changes in

access patterns [18]. Subsequently, the system determines whether the data belongs to cold or hot data based on a dynamic threshold *Threshold*, which is calculated as shown in equation (6).

$$\textit{Threshold} = t_e - \lfloor \log_{(1-\alpha)} St_k(t_n) \rfloor \tag{6}$$

In equation (6), t_e is the end time slice of the current block record. Traditional disaster recovery strategies often switch to any available backup node when the primary node fails, ignoring the differences between nodes, which may lead to a decrease in service quality after the switch. Therefore, this study introduces an intelligent backup node selection algorithm to select the node with the highest comprehensive utility from all available candidate nodes. The algorithm first identifies a list of cloud nodes that hold valid copies of the target data through blockchain metadata storage and evaluates the comprehensive service capability score Φ_m of each node. Then, the system uses a utility function $U(cl_m)$ to give each candidate node a final score and selects the node with the highest utility value as the final failover target [19]. This utility function is used to balance service quality and recovery costs, as shown in equation (7).

$$U(cl_m) = w_{ability} \cdot \Phi_m - w_{cost} \cdot C_{failover}(cl_m) \tag{7}$$

In equation (7), $C_{failover}(cl_m)$ is the estimated unit time operating cost after switching to the node. $w_{ability}$ and w_{cost} are weight coefficients that can be adjusted by administrators based on the characteristics of different businesses. The final selected node cl_{opt} is shown in equation (8).

$$cl_{opt} = \arg \max_{cl_m \in \text{Candidates}} U(cl_m) \tag{8}$$

Through equation (8), disaster recovery switching is no longer blind, but based on data-driven intelligent decision-making with the goal of achieving optimal business recovery results. To ensure the high reliability of the switching process itself and prevent recovery failures caused by single point failures or erroneous decisions in the control plane, this study introduces a recovery mechanism based on PBFT consensus algorithm. The process of PBFT algorithm is shown in Figure 5.

In Figure 5, after entering the consensus decision-making stage, the disaster recovery control node (which can be a highly available component in the management plane) initiates a proposal to switch to cl_{opt} to

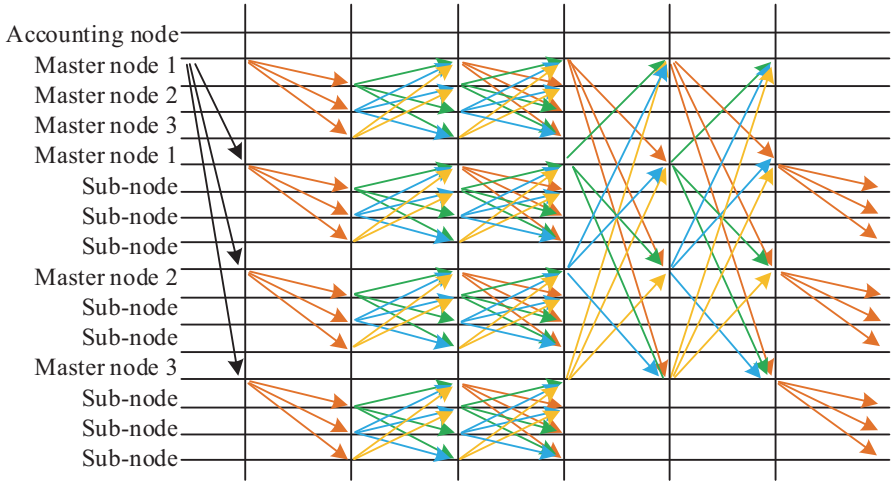

Figure 5 Algorithm flow of PBFT.

a set of predefined, distributed peer management nodes. These nodes will independently perform secondary confirmation on cl_{opt}'s health status and data validation results. Only when more than two-thirds of the nodes reply "agree", will the control node issue the final "submit switch" command [20]. Separately at the CDN layer, traffic redirection is optimized by clustering user request origins to minimize latency and cache-miss rates. An optimized K-Means with density-based seeding selects the highest-density points as initial centroids; the objective in (9) is then minimized over features such as client geo/ASN and real-time RTT/loss, as illustrated in Figure 6.

In Figure 6, the density and average density of data sample objects are calculated, and the data object with the highest density is selected as the initial clustering center. The specific objective function is shown in equation (9).

$$\max_{\pi} \mathrm{E}_{s,a\sim\pi}\left[\sum_{t=0}^{\infty}\gamma^{t}R(s_t,a_t)\right] \tag{9}$$

In equation (9), π is the scheduling strategy. $R(s_t, a_t)$ is the user access experience rating function. γ is the discount factor. The system continuously explores and iteratively updates node traffic allocation through feedback, thereby improving adaptive disaster recovery capabilities in uncertain environments. The proposed model is named HMM-IODR.

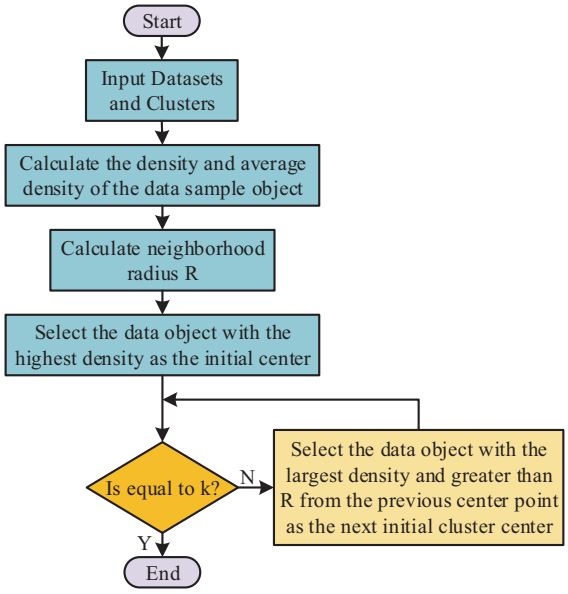

Figure 6 Optimized K-Means algorithm.

4 System Prototype Realization and Experimental Evaluation

4.1 System performance evaluation

The experiments were executed on a server equipped with an AMD EPYC 32-core CPU, 512 GB of memory, and two NVIDIA A100 GPUs, running Ubuntu 22.04. Models were implemented in TensorFlow 2.12. To rigorously evaluate the performance of our proposed HMM-IODR model, it was compared against four baseline models representing distinct approaches to multi-cloud management and disaster recovery. The first baseline, Integrated Multi-Cloud Network and Data Management Model (IMND-M), represents an advanced, integrated approach where control-plane orchestration and data-plane replication are tightly coupled, often using a consensus protocol to enforce strong consistency. In contrast, the Multi-Cloud Strategies for Enhanced Resilience and Flexibility (MS-ERF) model embodies a more traditional, rule-based portfolio where strategy selection is reactively triggered by SLO violations or budget constraints, typically without cross-layer co-optimization. A different perspective is offered by the Systemic Risk and Vulnerability Analysis of Multi-Cloud Environments (SRVA-MCE), which

focuses on risk assessment rather than active orchestration; this model analyzes service dependency graphs and uses stress tests to guide conservative, often manual, switching decisions. Lastly, the Disaster Recovery Planning with Multi-Cloud Replication (DRP-MCR) baseline reflects a conventional, data-centric disaster recovery planning approach, where the main focus is on configuring data replication and executing a predefined recovery plan to meet RPO and RTO targets. The dataset consists of public Network Diagnostic Tool (NDT) measurements curated by Measurement Lab (M-Lab), providing throughput and latency indicators via BigQuery. Data were split 80/20 for training/testing. The Adam optimizer was used with a learning rate of 0.001 for 150 epochs. Figure 7 reports the fitting performance of the compared models on the network disaster-recovery prediction task. That provider names appearing later (AWS, Azure, Google Cloud, Alibaba Cloud, Tencent Cloud) denote scenario labels on a virtual multi-cloud testbed rather than official provider-internal datasets[21]

Figure 7 shows the fitting effect of four multi cloud management models in service performance prediction. Figure 7(a) shows the HMM-IODR model, whose prediction results are highly consistent with the target value, with a coefficient of determination $R^2 = 0.97$, indicating that the model has extremely strong prediction accuracy. In Figure 7(b), the IMND-M model has a good fit but slightly lower accuracy, $R^2 = 0.91$. The prediction bias of MS-ERF in Figure 7(c) increases, $R^2 = 0.87$. The SRVA-MCE fitting in Figure 7(d) is the weakest, $R^2 = 0.79$, indicating that it is difficult to accurately reflect the state changes in complex cloud environments. The comparative analysis of different models in terms of the accuracy of master node selection and consensus completion time is shown in Figure 8.

Figure 8(a) shows that at different node sizes, the number of times abnormal nodes are selected as primary nodes in HMM-IODR is significantly less than in IMND-M. Especially when the number of nodes reaches 50, the error rate of HMM-IODR significantly increases slowly, indicating that it has stronger node quality perception and scheduling accuracy. Figure 8(b) shows the consensus completion time of the two models under different proportions of abnormal nodes (10% and 20%). HMM-IODR always maintains a shorter completion time in each round of consensus, especially in large-scale consensus (such as 5,000 rounds), where its performance advantage is more obvious, verifying its optimization ability for delay control and resource scheduling in disaster recovery switching. The performance evaluation results of different multi cloud management models in disaster recovery prediction scenarios are shown in Table 1.

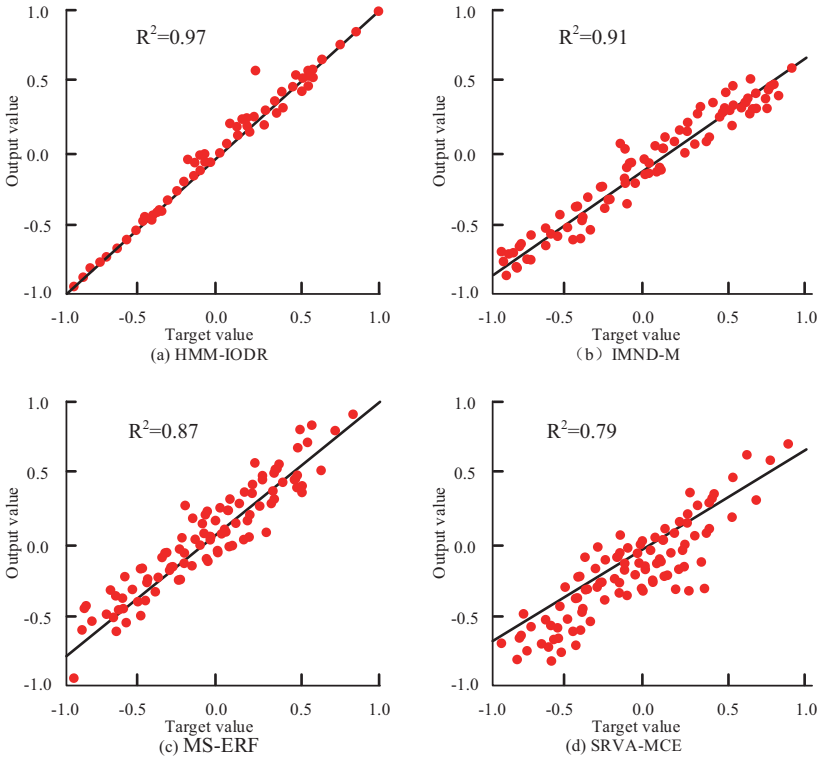

Figure 7 Fitting performance of multi-cloud management models on the network disaster-recovery prediction task using M-Lab NDT measurements (80/20 train/test split).

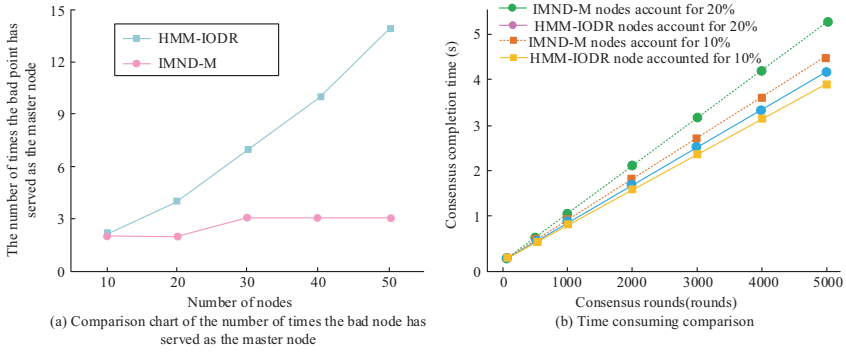

Figure 8 Comparison of model accuracy in master node selection and consensus completion time.

Table 1 System performance evaluation of four models in different scenarios

Model Name	Avg. Prediction Latency (ms)	Max Prediction Error (%)	Data Sync Time (s)	Resource Utilization (%)	Disaster Recovery Time (s)	Avg. Consensus Rounds	Model Convergence Time (min)
HMM-IODR	188.71	3.62	0.983	84.73	3.45	36.4	21.78
IMND-M	239.12	5.91	1.373	76.38	5.84	42.8	26.31
MS-ERF	274.58	6.88	1.621	72.45	6.73	48.5	28.92
SRVA-MCE	293.67	8.13	1.802	70.12	7.26	52.9	30.06

As shown in Table 1, HMM-IODR yielded quantitative advantages across all indicators. It achieved an average prediction delay of 188.71 ms and a maximum error of 3.62%, which were 21.1% and 38.7% lower than the next-ranked model, respectively. It also recorded a data synchronization time of under 1 second and a resource utilization rate of 84.73%. In addition, the model's disaster recovery switching time was 3.45 seconds, and its average of 36.4 consensus rounds and 21.78-minute convergence time were also lower than the other evaluated models. In contrast, IMND-M, MS-ERF, and SRVA-MCE all have varying degrees of disadvantages in terms of latency, synchronization efficiency, and resource scheduling. This validates the comprehensive adaptability and system performance of HMM-IODR in dealing with complex multi cloud environment disaster recovery tasks.

4.2 Performance evaluation of disaster recovery strategy

The experiment constructs a regional fault simulation environment on a virtual multi cloud testing platform, using a combination of five mainstream cloud service provider nodes: AWS, Azure, GCP, Alibaba Cloud, and Tencent Cloud. Each group of experiments deploys container-based microservice applications and simulates cross-cloud region unavailability events through fault injection tools. The experiment is conducted on a server cluster equipped with AMD EPYC 32 core CPU and 256 GB of memory. The controller node runs Ubuntu 22.04 and uses Prometheus to monitor latency and throughput. DNS failover is completed by CoreDNS in conjunction with a custom module on the testing platform. The data collection comes from real network performance samples provided by M-Lab (bandwidth, latency, connection loss, etc.), combined with self built service status log data, to evaluate the recovery latency and traffic stabilization process. Various disaster recovery strategies include Active-Passive, Cold Standby, Active-Active, zone transfer combined with DNS switching, and CDN intelligent redirection solutions.

Table 2 Performance comparison of different strategies in regional fault simulation

Experiment ID	Simulated Cloud Combination	Disaster Recovery Strategy	Failure Detection Delay (s)	Auto Failover Time (s)	RTO (s)	Through-put Drop (%)	Recovery Latency (ms)
Exp-01	AWS+Azure	Active-Passive	3.17	6.25	9.42	18.47	241.3
Exp-02	GCP+Azure	Cold Standby	4.83	25.11	29.94	34.26	356.9
Exp-03	AWS+GCP	Active-Active	2.03	1.42	3.45	9.83	188.7
Exp-04	AWS+ Alibaba Cloud	Regional Replication + DNS	3.76	12.88	16.64	22.11	278.5
Exp-05	GCP+ Tencent Cloud	CDN Redirection + Cold Standby	5.21	21.03	26.24	29.57	333.1

The strategy switching logic is coordinated and controlled through a unified disaster recovery controller. The performance comparison results of different disaster recovery strategies in regional fault simulation under multi-cloud architecture are shown in Table 2.

The experiment in Table 2 shows significant differences in the performance of each strategy. For example, the active passive strategy combination of AWS+Azure performs well on multiple indicators; The cold backup strategy combination of GCP+Azure performs poorly in indicators such as fault detection latency, automatic failover time, and RTO. In contrast, the active strategy combination of AWS+GCP performs well in all indicators, especially with an RTO of only 3.45 seconds, a throughput decrease of 9.83%, and a recovery delay of 188.7 milliseconds. The performance indicators of other combinations such as AWS+Alibaba Cloud and GCP+Tencent Cloud are between the two. This indicates that there are significant differences in fault recovery efficiency, throughput reduction, and latency when different Simulated Cloud Combinations are combined with disaster recovery strategies. The throughput and consensus latency performance comparison of multi cloud management models under different node sizes is shown in Figure 9.

Figure 9(a) shows that as the number of nodes increases from 40 to 100, the throughput of IMND-M continues to decline, dropping from 430 TPS to nearly 0, demonstrating poor scalability. HMM-IODR maintains a high throughput level under different replica factor settings (K = 4, 7, 10), with larger K values indicating more stable throughput and good scalability and resource tolerance. Figure 9(b) illustrates that during the same node addition process, the IMND-M consensus delay rapidly increases and eventually

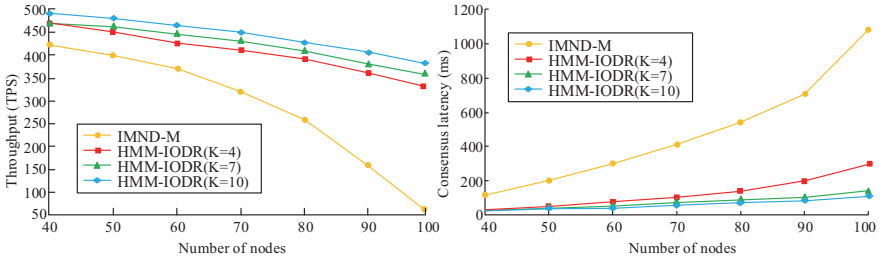

(a) Throughput comparison (b) Delay comparison

Figure 9 Comparison of model throughput and consensus latency under different node sizes.

exceeds 1,100 ms. HMM-IODR maintains low latency, especially under the setting of K = 10, which can be controlled within 120 ms even under the condition of 100 nodes. To further verify the adaptability of disaster recovery strategies in complex multi-cloud environments, experiments are conducted to evaluate the performance of five strategy combinations in multiple system operation dimensions (node selection accuracy, packet loss rate, network jitter, recovery success rate, memory usage, CPU load, etc.). The key performance indicators of different combinations of disaster recovery strategies in a multi-cloud environment are compared in Table 3.

Table 3 Comparison of key performance indicators

Experiment ID	Strategy Combination	Master Node Selection Accuracy (%)	Avg. Packet Loss Rate (%)	Network Jitter (ms)	Service Resumption Success Rate (%)	Max Memory Usage (MB)	Avg. CPU Load (%)
Exp-01	Active-Passive (AWS+Azure)	93.71	1.48	18.62	98.03	1583.4	62.87
Exp-02	Cold Standby (GCP+Azure)	86.35	2.27	26.94	91.42	1652.9	59.41
Exp-03	Active-Active (AWS+GCP)	97.82	0.93	13.25	99.67	1471.6	66.58
Exp-04	Regional Replication + DNS (AWS+Ali)	91.16	1.73	20.74	96.29	1538.2	61.02
Exp-05	CDN Redirection + Cold Standby (GCP+TC)	88.09	1.96	23.47	94.11	1604.3	60.35

As detailed in Table 3, the Active-Active combination (AWS+GCP) achieved the highest-ranking values across several key indicators. Its master node selection accuracy reached 97.82%, with a packet loss rate of 0.93% and a service recovery success rate of 99.67%, which indicates high stability and network tolerance. In contrast, the Cold Standby (GCP+Azure) strategy has weaker performance, especially in terms of network jitter and recovery success rate. Other strategies show intermediate levels in different dimensions, indicating that the effectiveness of multi-cloud disaster recovery solutions is closely related to Simulated Cloud Combination and strategy design, and needs to be flexibly selected based on business needs.

To provide direct guiding significance for industrial applications, this section presents a comparative experimental analysis of several key disaster recovery strategies, including Active-Active (multi-live deployment), Active-Passive (hot standby), Cold Standby, and intelligent CDN-based redirection. The reliability and real-world applicability of these findings are substantiated by leveraging the M-Lab Network Diagnostic Tool (NDT) public dataset for evaluation. This open dataset, comprising large-scale network performance measurements, ensures that the simulation of cross-cloud regional unavailability events and the assessment of recovery metrics are grounded in realistic conditions, thereby enhancing the credibility of the results. The experiment constructs a regional fault simulation environment on a virtual multi-cloud testing platform, using a combination of five mainstream cloud service provider nodes: AWS, Azure, GCP, Alibaba Cloud, and Tencent Cloud. Each experiment deploys container-based microservice applications, and fault injection tools simulate unavailability events. The experimental environment is a server cluster equipped with an AMD EPYC 32-core CPU and 256 GB of memory. The controller node runs Ubuntu 22.04 and uses Prometheus for monitoring, while DNS failover is managed by CoreDNS. The performance comparison results are summarized in Table 4.

As shown in Table 4, the proposed HMM-IODR model exhibits excellent performance in multi cloud disaster recovery scenarios, with a high fitting accuracy of $R^2 = 0.97$ for network disaster recovery prediction, an average prediction delay of only 188.71 milliseconds, and a maximum prediction error controlled at 3.62%. In the evaluation of specific disaster recovery strategies, the advantage of the Active Active strategy is most significant, achieving a recovery time objective (RTO) of 3.45 seconds in regional fault simulation, far exceeding other strategies such as cold backup (29.94 seconds) and hot backup (9.42 seconds). At the same time, the success rate of service recovery under this strategy reached 99.67%, the accuracy of primary

Table 4 Comparative performance analysis of key disaster recovery strategies in regional fault simulation

Experiment ID	Strategy Combination	RTO (s)	Throughput Drop (%)	Service Resumption Success Rate (%)	Avg. CPU Load (%)
Exp-03	Active-Active (AWS+GCP)	3.45	9.83	99.67	66.58
Exp-01	Active-Passive (AWS+Azure)	9.42	18.47	98.03	62.87
Exp-05	CDN Redirection + Cold Standby (GCP+TC)	26.24	29.57	94.11	60.35
Exp-02	Cold Standby (GCP+Azure)	29.94	34.26	91.42	59.41

node selection was 97.82%, and the throughput decrease was only 9.83%, demonstrating extremely high reliability and business continuity assurance capabilities. In terms of scalability, when the number of nodes increases to 100, the HMM-IODR model can still maintain consensus latency within 120 milliseconds and maintain high throughput levels, significantly better than the baseline model whose performance deteriorates rapidly with increasing nodes.

4.3 Economic cost analysis

While the performance evaluation in Section 4.2 demonstrates the superior recovery time and stability of the Active-Active strategy, a comprehensive assessment for industrial implementation must also consider the Total Cost of Ownership (TCO). A decision based solely on performance metrics risks a "performance-cost" imbalance, as the most resilient architectures are often the most expensive. To address this, a comparative economic analysis is presented based on typical cloud provider pricing models. The TCO for a disaster recovery strategy can be deconstructed into several key components: compute resource costs, storage and data transfer fees, and management overhead. The Active-Active (multi-live) strategy inherently incurs higher costs, as it requires maintaining two or more fully provisioned, production-grade environments running in parallel. This effectively doubles the expenditure on compute instances, databases, and other core services. In contrast, the Active-Passive (hot standby) model allows the backup site to operate with scaled-down resources that are only fully provisioned upon failover, leading to significantly lower baseline compute costs. Furthermore, cross-region or cross-cloud data synchronization generates substantial data egress charges. The continuous, often bi-directional, replication in an Active-Active setup

typically results in higher data transfer volumes than the one-way synchronization characteristic of an Active-Passive architecture. Finally, the operational complexity of managing a global traffic manager and ensuring consistency across multiple live sites increases the management overhead for the Active-Active strategy. Therefore, the choice of strategy represents a critical trade-off: while the Active-Active approach offers the lowest RTO (3.45 s) and minimal performance degradation, its TCO can be substantially higher than that of an Active-Passive solution. The optimal choice is contingent on business requirements; mission-critical applications where downtime incurs severe financial or reputational damage may justify the expense of an Active-Active architecture, whereas more cost-sensitive services might find the balance of cost and resilience offered by a hot standby strategy to be more appropriate.

5 Conclusion

As multi-cloud deployment increasingly becomes the standard for enterprises supporting critical business operations, designing a management architecture that ensures high availability and intelligent disaster recovery has become an urgent challenge. This study proposes the HMM-IODR model, which integrates a blockchain-based verification mechanism, an active/standby dynamic scheduling algorithm, the Practical Byzantine Fault Tolerance (PBFT) protocol, and CDN-level intelligent traffic allocation. The model introduces innovative designs in both system architecture and disaster recovery strategies.

To evaluate its effectiveness, a real-world virtual multi-cloud testbed was constructed to simulate cross-cloud regional fault scenarios. Performance was assessed using a combination of Measurement Lab's NDT network performance data and self-generated log samples. Provider names are used solely to represent testbed scenarios, without implying provider-specific claims.

Experimental results demonstrate that HMM-IODR achieves high prediction accuracy in network disaster recovery, with a fitting accuracy of $R^2 = 0.97$, a prediction delay of 188.71 ms, and a maximum error of 3.62%. The consensus completion time significantly outperformed the baseline model, requiring no more than 610 ms under 5,000 consensus attempts. In regional fault recovery simulations, the HMM-IODR model combined with the Active–Active strategy yielded a recovery time objective (RTO) of 3.45 s, a recovery delay of 188.7 ms, and throughput degradation of less than 10%. The master node selection achieved 97.82% accuracy, the recovery

success rate reached 99.67%, and resource utilization was efficient, with CPU utilization at 66.58% and memory consumption at 1471.6 MB.

While the proposed approach demonstrates stability and high performance across multiple metrics, the current system primarily focuses on optimizing fundamental network and service status data, with limited attention to multi-tenant security isolation. Future work could therefore focus on addressing this gap by integrating a zero-trust architecture to enhance security. Additionally, incorporating intelligent algorithms such as federated learning and reinforcement learning could build a self-learning disaster recovery control center, thereby enhancing adaptability and robustness in responding to previously unseen fault scenarios.

References

[1] Ullah A, Kiss T, Kovács J, et al. Orchestration in the Cloud-to-Things compute continuum: taxonomy, survey and future directions[J]. Journal of Cloud Computing, 2023, 12(1): 1–29. DOI: https://doi.org/10.1186/s13677-023-00516-5.

[2] Ouchaou L, Nacer H, Labba C. Towards a distributed SaaS management system in a multi-cloud environment[J]. Cluster Computing, 2022, 25(6): 4051–4071. DOI: https://doi.org/10.1007/s10586-022-03619-x.

[3] Lefranc G, Lopez-Juarez I, Gatica G. Enhancing FMS Performance through Multi-Agent Systems in the Context of Industry 4.0[J]. Studies in Informatics and Control, 2024, 33(2): 5–14. DOI: https://doi.org/10.24846/v33i2y202401.

[4] Sun X, Chen J, Zhao H, Zhang W, Zhang Y. Sequential disaster recovery strategy for resilient distribution network based on cyber–physical collaborative optimization[J]. IEEE Transactions on Smart Grid, 2022, 14(2): 1173–1187. DOI: https://doi.org/10.1109/TSG.2022.3198696.

[5] Iorio M, Risso F, Palesandro A, Camiciotti L, Manzalini A. Computing without borders: The way towards liquid computing[J]. IEEE Transactions on Cloud Computing, 2022, 11(3): 2820–2838. DOI: https://doi.org/10.1109/TCC.2022.3229163.

[6] Li Y, Hwang K, Shuai K, Li Z, Zomaya A. Federated clouds for efficient multitasking in distributed artificial intelligence applications[J]. IEEE Transactions on Cloud Computing, 2022, 11(2): 2084–2095. DOI: https://doi.org/10.1109/TCC.2022.3184157.

[7] Zhang FL. Evolutionary Algorithm for Dynamic Resource Allocation and Its Applications[J]. International Journal of Simulation Modelling, 2024, 23(3): 531–542. DOI: https://doi.org/10.2507/IJSIMM23-3-CO14.

[8] Addya SK, Satpathy A, Ghosh BC, Chakraborty S, Ghosh SK, Das SK. CoMCLOUD: Virtual machine coalition for multi-tier applications over multi-cloud environments[J]. IEEE Transactions on Cloud Computing, 2023, 11(1): 956–970. DOI: https://doi.org/10.1109/TCC.2021.3122445.

[9] Globa L, Kartashov A. Optimizing distributed data storage in multi-cloud environments: algorithmic approach[J]. Information and Telecommunication Sciences, 2024, (2): 4–12. DOI: https://doi.org/10.20535/2411-2976.22024.4-12.

[10] Zhang T, Liu C, Tian Q, Cheng B. Cloud-Edge Collaboration-Based Multi-Cluster System for Space-Ground Integrated Network[J]. International Journal of Satellite Communications and Networking, 2025, 43(1): 40–60. DOI: https://doi.org/10.1002/sat.1541.

[11] Kim B, Calin D, Tenny N, Shariat M, Fan M. Device centric distributed compute, orchestration and networking[J]. IEEE Wireless Communications, 2023, 30(4): 6–8. DOI: https://doi.org/10.1109/MWC.2023.10251878.

[12] Benmerar TZ, Theodoropoulos T, Fevereiro D, Rosa L, Rodrigues J, Taleb T, Barone P, Giuliani G, Tserpes K, Cordeiro L. Towards establishing intelligent multi-domain edge orchestration for highly distributed immersive services: a virtual touring use case[J]. Cluster Computing, 2024, 27(4): 4223–4253. DOI: https://doi.org/10.1007/s10586-024-04413-7.

[13] Taghinezhad-Niar A, Taheri J. Reliability, rental-cost and energy-aware multi-workflow scheduling on multi-cloud systems[J]. IEEE Transactions on Cloud Computing, 2023, 11(3): 2681–2692. DOI: https://doi.org/10.1007/s10586-024-04413-710.1109/TCC.2022.3223869.

[14] Tusa F, Clayman S. End-to-end slices to orchestrate resources and services in the cloud-to-edge continuum[J]. Future Generation Computer Systems, 2023, 141: 473–488. DOI: https://doi.org/10.1016/j.future.2022.11.026.

[15] Hegyi P. Service deployment design in latency-critical multi-cloud environment[J]. Computer Networks, 2022, 213: 108975. DOI: 10.1016/j.comnet.2022.108975.

[16] Ashrafi R, AlKindi H. A framework for IS/IT disaster recovery planning[J]. International Journal of Business Continuity and Risk Management, 2022, 12(1): 1–21. DOI:10.1504/IJBCRM.2022.10045649.
[17] Mišić J, Mišić VB, Chang X. Design of proof-of-stake PBFT algorithm for IoT environments[J]. IEEE Transactions on Vehicular Technology, 2022, 72(2): 2497–2510. DOI:10.1109/TVT.2022.3213226.
[18] Luo H, Yang X, Yu H, Sun G, Lei B, Guizani M. Performance analysis and comparison of nonideal wireless PBFT and RAFT consensus networks in 6G communications[J]. IEEE Internet of Things Journal, 2023, 11(6): 9752–9765. DOI: 10.1109/JIOT.2023.3323492.
[19] Kontodimas K, Soumplis P, Kretsis A, Kokkinos P, Fehér M, Lucani DE, Varvarigos E. Secure distributed storage orchestration on heterogeneous cloud-edge infrastructures[J]. IEEE Transactions on Cloud Computing, 2023, 11(4): 3407–3425. DOI: 10.1109/TCC.2023.3287653.
[20] Castro M, Liskov B. Practical byzantine fault tolerance and proactive recovery[J]. ACM Transactions on Computer Systems (TOCS), 2002, 20(4): 398–461. DOI: 10.1145/571637.571640.
[21] Measurement Lab (M-Lab). M-Lab NDT Datasets (BigQuery) [DB/OL]. Available: https://www.measurementlab.net/data/. (Accessed 2025-08-11).

Biographies

Wenchao Li (August 1981–), male, graduated from Hohai University with a master's degree in Power System and Automation. After graduation, I worked as a senior engineer at the Power Dispatch Control Center of China Southern Power Grid Co., Ltd. My current research direction is engaged in power system automation work.

Guang Ma (September 1995–), male, graduated from Zhejiang University with a master's degree in Power System and Automation. After graduation, I worked as an engineer at the Power Dispatch Control Center of China Southern Power Grid Co., Ltd. My current research direction is engaged in power system automation work.

Wenchong Fang (January 1986–), male, graduated from Sun Yat sen University with a master's degree in Computer Application Technology. After graduation, I worked as a senior engineer at the Power Dispatch Control Center of China Southern Power Grid Co., Ltd. My current research direction is engaged in power system automation work.

Xiqi He (April 1983–), male, graduated from North China Electric Power University with a master's degree in Power System and Automation. After graduation, I worked as a senior engineer at the Power Dispatch Control Center of China Southern Power Grid Co., Ltd. My current research direction is engaged in power system automation work.

Jin Li (May 1979–), male, graduated from Central South University with a master's degree in Power System Automation. After graduation, I worked as a professor level senior engineer at the Power Dispatch Control Center of China Southern Power Grid Co., Ltd. My current research direction is engaged in power system automation work.