Geo-Distributed Quantum Hyper-Aether BC/DR: A Semantic-Aware and Knowledge-Resilient Disaster Recovery Architecture

Hirofumi Inomata

Abstract

This paper presents Geo-Distributed Quantum Hyper-Aether BC/DR, an AI-native business continuity and disaster recovery architecture for wide-area distributed systems. Conventional BC/DR mechanisms mainly focus on workload relocation, infrastructure redundancy, and data replication. However, AI-native systems require additional recovery dimensions, including semantic coverage, knowledge preservation, VM lifecycle reconstruction, and adaptation to changing semantic demand. The proposed architecture treats semantic tasks, virtual machine lifecycle events, and knowledge recovery as first-class control units across multiple regions.

A final simulation model is constructed using a three-region environment, four recovery methods, and five wide-area disaster scenarios. The compared methods are Static Replication, Load-based Failover, Semantic Routing, and the proposed Geo-Distributed QHA BC/DR. The disaster scenarios include single region failure, inter-region network partition, knowledge repository loss, sudden semantic demand shift, and multi-region cascading failure. The evaluation metrics are recovery time, semantic coverage, knowledge loss, service continuity, and transfer cost.

Simulation results show that the proposed Geo-QHA BC/DR method achieves the best overall performance. It provides the shortest recovery time, highest semantic coverage, lowest knowledge loss, highest service continuity, and lowest transfer cost among the compared methods. The advantage is particularly clear in knowledge repository loss, sudden semantic demand shift, and multi-region cascading failure scenarios. These results suggest that future AI-native BC/DR systems should not only replicate infrastructure and data, but also preserve meaning, reconstruct knowledge, reuse dormant VMs, and autonomously reconfigure distributed semantic control.

1 Introduction

Business continuity and disaster recovery, abbreviated as BC/DR, are essential for maintaining service availability under failures. Traditional BC/DR systems are usually designed around infrastructure-level redundancy, storage replication, workload migration, backup restoration, and network failover. These approaches are effective when the primary objective is to restart the same workload in another site.

However, AI-native distributed systems introduce a different recovery problem. In such systems, service continuity depends not only on whether servers and data are available, but also on whether the system still preserves the semantic knowledge required to answer user requests, select appropriate reasoning paths, and adapt to changing demand. If a disaster destroys or isolates a knowledge repository, a conventional recovery method may restore compute resources but still fail to recover the semantic capability of the system.

Quantum Hyper-Aether, abbreviated as QHA, is a conceptual AI-native distributed computing substrate in which virtual machines, language models, neural networks, semantic tasks, purpose, context, and knowledge are treated as dynamic system elements. In this paper, the QHA concept is extended to a geo-distributed BC/DR architecture. The proposed approach, called Geo-Distributed QHA BC/DR, coordinates VM lifecycle control, semantic routing, dormant VM reuse, and autonomous knowledge recollection across multiple regions.

The main contribution of this paper is a final simulation-based evaluation of the proposed architecture. The simulation compares four BC/DR methods under five wide-area disaster scenarios and evaluates them using five metrics: recovery time, semantic coverage, knowledge loss, service continuity, and transfer cost.

2 Background and Motivation

2.1 Limitations of Infrastructure-Centered BC/DR

Conventional BC/DR methods typically assume that recovery means restoring infrastructure, network connectivity, storage consistency, and workload execution. Static replication maintains copies of workloads and data in multiple sites. Load-based failover redirects requests to less damaged or less loaded regions. These mechanisms are important, but they are not sufficient for AI-native systems.

In AI-native systems, the operational state includes learned knowledge, task context, semantic embeddings, routing policies, reasoning histories, and domain-specific constraints. Therefore, even if virtual machines are restarted, the service may still lose the ability to respond to specific semantic tasks. This motivates the need for semantic-aware and knowledge-resilient recovery.

2.2 Semantic Recovery Requirement

A semantic task is a request or computational unit whose meaning determines how it should be routed, processed, or recovered. For example, two requests may have similar computational load but completely different required knowledge domains. A conventional load balancer sees them as similar, whereas a semantic BC/DR controller should treat them differently.

In this paper, semantic recovery means preserving or reconstructing the ability to process meaning-specific tasks after a disaster. This includes the following capabilities:

- maintaining semantic coverage across regions,
- reducing loss of domain knowledge,
- detecting missing or degraded knowledge,
- reusing dormant VMs that retain useful knowledge,
- recollecting knowledge when deleted or failed VMs cause capability loss,
- dynamically reassigning semantic tasks to appropriate regions.

3 Proposed Architecture

3.1 Overview

The proposed Geo-Distributed QHA BC/DR architecture consists of three regional sites connected by a semantic control layer. Each region contains VMs, local knowledge stores, semantic task queues, and routing controllers. The global QHA control layer observes disaster conditions and reconfigures the system using semantic information.

The architecture has the following major components:

- **Semantic Task Layer**: represents incoming tasks by meaning, context, and required knowledge.
- **VM Lifecycle Layer**: controls VM creation, promotion, dormancy, reuse, merge, rejection, and deletion.
- **Knowledge Recovery Layer**: detects knowledge loss and recollects or reconstructs missing knowledge.

- **Geo-Distributed Routing Layer**: assigns tasks to regions according to semantic suitability, service continuity, and disaster state.
- **BC/DR Control Layer**: coordinates failover, recovery, and adaptation across regions.

3.2 VM Lifecycle Control

In the proposed architecture, a VM is not merely a compute instance. It is treated as a semantic processing unit with a purpose, internal knowledge, and task history. The lifecycle of a VM includes the following states:

- **Active**: the VM currently processes semantic tasks.
- **Dormant**: the VM is not actively processing tasks but retains useful knowledge.
- **Promoted**: a discovered or temporary VM is accepted as a stable VM.
- **Merged**: knowledge or function is integrated into another VM.
- **Rejected**: the VM is determined to be irrelevant or unsafe.
- **Deleted**: the VM is removed, possibly causing knowledge loss.

Dormant VMs are particularly important for BC/DR. During a disaster, dormant VMs may be reused to restore semantic coverage without fully rebuilding knowledge from scratch.

3.3 Knowledge Recollection

Knowledge recollection is triggered when the system detects that user tasks can no longer be answered with sufficient semantic coverage. This may occur after VM deletion, regional failure, repository loss, or sudden demand shift. The system then attempts to recollect or reconstruct knowledge from surviving regions, dormant VMs, external sources, or previously stored semantic summaries.

This mechanism differentiates Geo-QHA BC/DR from conventional replication. Instead of only restoring files or instances, the system attempts to restore the semantic capability required for future tasks.

4 Simulation Model

4.1 Compared Methods

The simulation compares four methods.

1. **Static Replication**: workloads and data are replicated, but semantic adaptation is limited.
2. **Load-based Failover**: requests are redirected according to region load and availability.
3. **Semantic Routing**: requests are routed using semantic suitability, but VM lifecycle and knowledge recovery are limited.
4. **Proposed Geo-QHA BC/DR**: semantic routing, VM lifecycle control, dormant VM reuse, and knowledge recollection are jointly applied.

4.2 Disaster Scenarios

Five wide-area disaster scenarios are modeled.

1. **Single Region Failure**: one region becomes unavailable.
2. **Inter-region Network Partition**: communication among regions is partially disrupted.
3. **Knowledge Repository Loss**: a major knowledge repository is lost or corrupted.
4. **Sudden Semantic Demand Shift**: the distribution of user tasks changes rapidly.
5. **Multi-region Cascading Failure**: multiple regions and links degrade sequentially.

4.3 Evaluation Metrics

The following five metrics are used.

- **Recovery Time**: average time required to restore useful service.
- **Semantic Coverage**: fraction of required semantic capability preserved or recovered.
- **Knowledge Loss**: fraction of useful knowledge lost after recovery.
- **Service Continuity**: degree to which service remains available during failure.
- **Transfer Cost**: relative cost of inter-region data, task, and knowledge movement.

4.4 Abstract Risk Model

Each scenario is represented by five damage factors:

$$D = (d_{\text{infra}}, d_{\text{network}}, d_{\text{knowledge}}, d_{\text{semantic}}, d_{\text{cascade}})$$

where each component ranges from 0 to 1. Each method has a corresponding capability vector:

$$C = (c_{\text{infra}}, c_{\text{network}}, c_{\text{knowledge}}, c_{\text{semantic}}, c_{\text{cascade}}, c_{\text{transfer}})$$

The residual damage for a factor is defined as:

$$R_i = D_i(1 - C_i)$$

The total risk is modeled as a weighted combination:

$$R = 0.25R_{\text{infra}} + 0.20R_{\text{network}} + 0.25R_{\text{knowledge}} + 0.20R_{\text{semantic}} + 0.10R_{\text{cascade}} + \epsilon$$

where ϵ is random trial noise.
Recovery time is modeled as:

$$T_{\text{rec}} = 10 + 95R + 20R_{\text{network}} + 18R_{\text{cascade}}$$

Semantic coverage is modeled as:

$$S_{\text{cov}} = 1 - 0.48R_{\text{semantic}} - 0.38R_{\text{knowledge}} - 0.16R_{\text{network}} - \epsilon_s$$

Knowledge loss is modeled as:

$$K_{\text{loss}} = 0.70R_{\text{knowledge}} + 0.18R_{\text{semantic}} + 0.12R_{\text{cascade}} + \epsilon_k$$

Service continuity is modeled as:

$$C_{\text{svc}} = 0.98 - 0.55R - 0.0025T_{\text{rec}} + 0.18S_{\text{cov}} - \epsilon_c$$

Transfer cost is modeled as:

$$C_{\text{transfer}} = 35 + 105(d_{\text{network}} + 0.6d_{\text{cascade}})(1 - c_{\text{transfer}}) + 35R_{\text{knowledge}} + \epsilon_t$$

All bounded metrics are clipped to the interval $[0, 1]$ where appropriate.

5 Results

5.1 Overall Ranking

The simulation result is summarized in Table 1. Higher values are better for semantic coverage, service continuity, and composite score. Lower values are better for recovery time, knowledge loss, and transfer cost.

Table 1: Overall simulation results.

Method	Rec. Time	Sem. Cov.	Know. Loss	Svc. Cont.	Trans. Cost	Score
Proposed Geo-QHA BC/DR	22.98	0.918	0.077	0.993	56.16	1.000
Semantic Routing	34.44	0.819	0.193	0.934	70.85	0.570
Load-based Failover	42.32	0.690	0.287	0.860	79.79	0.169
Static Replication	47.24	0.654	0.309	0.824	86.80	0.000

The proposed method achieves the best result in all five primary metrics. Compared with Static Replication, the proposed method reduces average recovery time from 47.24 to 22.98 and reduces knowledge loss from 0.309 to 0.077. It also increases semantic coverage from 0.654 to 0.918.

5.2 Scenario-Level Recovery Time

Table 2 shows recovery time under the five disaster scenarios.

Table 2: Recovery time across disaster scenarios. Lower is better.

Scenario	Static	Load-based	Semantic	Proposed
Single Region Failure	31.4	27.8	25.1	17.8
Inter-region Network Partition	47.1	42.1	35.6	24.9
Knowledge Repository Loss	42.8	39.3	31.0	20.6
Sudden Semantic Demand Shift	39.2	35.0	25.5	17.9
Multi-region Cascading Failure	75.7	67.4	55.1	33.7

The proposed method achieves the lowest recovery time in every scenario. The largest gains appear in severe distributed failures, especially multi-region cascading failure. This suggests that semantic-aware VM lifecycle control and knowledge recollection are effective when infrastructure, network, and knowledge failures occur together.

5.3 Interpretation

The results indicate three important findings.

First, semantic routing alone improves performance compared with static replication and load-based failover. This means that treating task meaning as a routing factor is useful even without full QHA lifecycle control.

Second, the proposed Geo-QHA BC/DR method outperforms semantic routing because it includes additional mechanisms for knowledge recovery and VM lifecycle adaptation. This is particularly important when the failure is not only physical but also semantic.

Third, knowledge loss is a central metric for AI-native BC/DR. Conventional methods may restore infrastructure but still lose the ability to answer domain-specific tasks. The proposed method reduces this risk by using dormant VMs and knowledge recollection.

6 Discussion

6.1 Why Geo-QHA Improves BC/DR

The proposed architecture improves BC/DR for three reasons.

- **Semantic-aware routing**: tasks are routed according to meaning and required knowledge, not only load or availability.
- **VM lifecycle resilience**: VMs can be reused, promoted, merged, or kept dormant to preserve semantic capability.
- **Knowledge recollection**: lost knowledge can be detected and reconstructed after VM deletion or regional failure.

These mechanisms allow the system to recover not only execution capacity but also semantic service capability.

6.2 Difference from Conventional DR

In conventional DR, a successful recovery often means that the system is restarted and data is restored. In AI-native systems, this criterion is incomplete. A system may be running but still unable to answer important questions if the relevant knowledge or semantic routing path has been lost.

Geo-QHA BC/DR expands the recovery target from infrastructure state to semantic state. The recovery objective becomes:

$$\text{Recover service} = \text{Recover compute} + \text{Recover data} + \text{Recover meaning} + \text{Recover knowledge}$$

This broader recovery target is the main conceptual contribution of the proposed architecture.

6.3 Implications for AI-Native Distributed Systems

The simulation suggests that future AI-native distributed systems should include semantic-aware BC/DR control as a native function. This may be particularly important for systems using multiple language models, distributed knowledge bases, autonomous agents, or dynamically generated VMs.

A practical implementation may require:

- semantic task embeddings,
- region-level knowledge inventories,
- VM purpose descriptors,
- dormant VM storage policies,

- knowledge loss detection,
- safe knowledge recollection mechanisms,
- human audit for high-risk knowledge reconstruction.

7 Limitations

This study is based on an abstract simulation model. The results should be interpreted as a conceptual and comparative evaluation, not as measurements from a deployed production system.

The main limitations are as follows.

- The scenario parameters are manually defined.
- The method capability values are assumed rather than measured.
- The simulation does not yet use real cloud failure traces.
- The semantic coverage metric is abstract and should be validated using real task datasets.
- Transfer cost is modeled as a relative cost, not as an actual network billing value.
- Security, consistency, privacy, and compliance constraints are not fully modeled.

Future work should evaluate the proposed architecture using real distributed workloads, real embeddings, realistic network failure traces, and measured recovery procedures.

8 Conclusion

This paper proposed and evaluated Geo-Distributed Quantum Hyper-Aether BC/DR, a semantic-aware and knowledge-resilient disaster recovery architecture for AI-native distributed systems. The proposed method treats semantic tasks, VM lifecycle events, dormant VM reuse, and knowledge recollection as first-class BC/DR control elements.

A final simulation using three regions, four recovery methods, and five wide-area disaster scenarios showed that the proposed method achieved the best overall performance. It reduced recovery time and knowledge loss while improving semantic coverage and service continuity. The results indicate that AI-native BC/DR should move beyond infrastructure replication and workload relocation toward semantic and knowledge-aware recovery.

Geo-Distributed QHA BC/DR therefore provides a promising framework for future business continuity and disaster recovery in AI-native computing environments where meaning, knowledge, context, and VM lifecycle dynamics are essential to service resilience.

References

[1] A. Wood, "Data Center Business Continuity and Disaster Recovery," *IT Professional*, vol. 5, no. 4, pp. 58–62, 2003.

[2] M. Alhazmi and Y. Malaiya, "Evaluating Disaster Recovery Plans Using the Cloud," *Proceedings of the Annual Reliability and Maintainability Symposium*, 2013.

[3] T. Berners-Lee, J. Hendler, and O. Lassila, "The Semantic Web," *Scientific American*, 2001.

[4] A. S. Tanenbaum and M. van Steen, *Distributed Systems: Principles and Paradigms*. Pearson, 2007.

[5] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. Yu, "A Comprehensive Survey on Graph Neural Networks," *IEEE Transactions on Neural Networks and Learning Systems*, 2021.

[6] W. X. Zhao et al., "A Survey of Large Language Models," *arXiv preprint arXiv:2303.18223*, 2023.