International Journal of Scientific Research and Engineering Development-– Volume 9 Issue 6, Nov – Dec 2026
Available at www.ijsred.com

RESEARCH ARTICLE OPEN ACCESS

Enterprise AI Disaster Recovery and Operational Resilience Architecture

A Comprehensive Enterprise Architecture for AI Continuity, Failover, Recovery, Degraded Operation, Dependency Resilience, and Business Service Restoration

Sanjeeve Kumar Gajadi*
*(HCL Technologies, Bengaluru, Karnataka, India)
Email: sanjeevekumar.g@hcltech.com

****

Abstract:
Enterprise AI services increasingly support customer interactions, operations, planning, analytics, and autonomous processes. Disruption of models, providers, vector stores, data platforms, GPUs, APIs, identity services, or agent tools can therefore create significant business impact. Traditional disaster recovery plans may not address AI-specific dependencies such as model endpoints, prompt repositories, retrieval indexes, evaluation baselines, policy services, and external foundation-model providers. This paper proposes an Enterprise AI Disaster Recovery and Operational Resilience Architecture that defines recovery objectives, dependency mapping, multi-region design, provider fallback, model portability, data recovery, configuration backup, degraded operation, manual fallback, resilience testing, and recovery governance. The architecture supports predictive AI, Generative AI, RAG, and agentic systems. The proposed approach enables enterprises to restore AI-enabled services safely while maintaining data integrity, security, governance, and business continuity.

Keywords — AI Disaster Recovery, Operational Resilience, Business Continuity, Generative AI, Agentic AI, RAG, Multi-Region Architecture, Model Failover, Enterprise Architecture, AI Governance

****

I. INTRODUCTION

AI services depend on a wider set of components than conventional applications and may fail even when core application infrastructure remains available.

Resilience architecture must identify the dependencies required to restore safe and useful AI behavior.

II. RESILIENCE OBJECTIVES

Recovery objectives should define acceptable downtime, data loss, model degradation, autonomy reduction, and manual intervention for each use case.

Critical AI functions should align with business continuity priorities rather than technical preferences alone.

III. DEPENDENCY MAPPING

Dependency maps should include models, providers, data sources, vector stores, prompts, tools, identity, policy engines, integration, observability, and network services.

Hidden external dependencies can become single points of failure if they are not included in resilience planning.

IV. RECOVERY STRATEGIES

Strategies may include multi-region deployment, secondary providers, alternate models, cached responses, read-only modes, or deterministic fallback workflows.

Recovery design should consider semantic and behavioral compatibility, not only infrastructure availability.

V. DATA AND KNOWLEDGE RECOVERY

Data products, vector indexes, embeddings, prompt repositories, model metadata, and policy configurations require backup and recovery procedures.

 ©IJSRED: All Rights are Reserved

Restored knowledge should be validated for freshness, completeness, permissions, and consistency before use.

VI. MODEL AND PROVIDER FAILOVER

Provider failover requires pretested alternative models, routing rules, performance thresholds, and security approvals.

Substitute models should not automatically inherit production approval if their behavior materially differs.

VII. AGENTIC AI SAFE DEGRADATION

Autonomous systems should support reduced-authority modes during degraded operation.

Agents may be limited to read-only actions, restricted tools, lower transaction limits, or mandatory human approval until full assurance is restored.

VIII. RECOVERY VALIDATION

Recovery tests should verify functionality, model quality, data access, security, policy enforcement, observability, and business outcomes.

A technically available service should not be declared recovered if critical governance or safety controls remain unavailable.

IX. RESILIENCE TESTING

Exercises should simulate provider outages, region failures, vector-store loss, identity disruption, policy-engine failure, and corrupted configuration.

Chaos testing and tabletop exercises can identify weak dependencies before real incidents occur.

X. OPERATIONAL GOVERNANCE

Resilience plans should define owners, invocation authority, communication, risk acceptance, and recovery evidence.

Changes to architecture, providers, models, and critical dependencies should trigger resilience-plan review.

XI. METRICS

Metrics include recovery time, recovery point, failover success, degraded-service duration, manual fallback effectiveness, data restoration completeness, and unresolved resilience risks.

Trend analysis should identify repeated recovery weaknesses and concentration risks.

XII. CONCLUSION

Enterprise AI Disaster Recovery and Operational Resilience Architecture extends continuity planning to AI-specific dependencies and behaviors.

By combining dependency visibility, tested fallback, safe degradation, and governance, organizations can restore AI-enabled business services with greater confidence.

ACKNOWLEDGMENT

The author acknowledges the contributions of enterprise architecture practitioners, AI governance specialists, standards bodies, security professionals, data leaders, and technology researchers whose published work informed the concepts synthesized in this paper.

REFERENCES

[1] National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework (AI RMF 1.0), 2023.
[2] National Institute of Standards and Technology, AI RMF Generative Artificial Intelligence Profile, 2024.
[3] ISO/IEC 42001:2023, Information technology — Artificial intelligence — Management system.
[4] ISO 31000:2018, Risk management — Guidelines.
[5] ISO/IEC 27001:2022, Information security, cybersecurity and privacy protection — Information security management systems.
[6] The Open Group, TOGAF Standard, 10th Edition, 2022.
[7] ISACA, COBIT 2019 Framework: Governance and Management Objectives, 2019.
[8] OWASP Foundation, OWASP Top 10 for Large Language Model Applications.
[9] MITRE, ATLAS: Adversarial Threat Landscape for Artificial-Intelligence Systems.
[10] OECD, Recommendation of the Council on Artificial Intelligence.
[11] SAP SE, SAP Business Technology Platform Documentation.
[12] SAP SE, SAP Business Data Cloud and SAP Datasphere Documentation.

 ©IJSRED: All Rights are Reserved