CHAPTER 1
INTRODUCTION

1.1 Background

In today's rapidly evolving digital landscape, businesses are increasingly reliant on cloud computing for operational resilience and business continuity. Cloud-based Disaster Recovery (DR) has emerged as a crucial solution for protecting data, ensuring service availability, and minimizing downtime during unexpected disruptions. This trend is particularly significant for Small and Medium Enterprises (SMEs), which often lack resources for on-premises disaster recovery and face threats such as cyberattacks, data breaches, hardware failures, and natural disasters, all of which can disrupt operations and cause significant data loss and financial damage. For example, a survey by the Disaster Recovery Journal (DRJ) found that hardware and system failures account for 49% of data loss incidents, followed by human errors at 36%, computer viruses at 7%, and natural disasters at 3% [1]. The increased frequency of incidents underscores the necessity for effective disaster recovery systems that ensure minimal downtime and data loss [2].

Among the cloud service providers, Amazon Web Services (AWS) and Microsoft Azure dominate the global market due to their feature rich platforms, global infrastructure, and extensive disaster recovery service portfolios. AWS offers services such as Elastic Disaster Recovery (EDR), Amazon S3 Cross-Region Replication, and AWS Backup, which help ensure business continuity with minimal Recovery Time Objectives (RTO) and Recovery Point Objectives (RPO) [3]. Similarly, Azure's offerings-including Azure Site Recovery (ASR) and Azure Backup-provide automated failover and data replication capabilities across multiple cloud regions [4].

Despite these robust offerings, relatively little research directly compares AWS and Azure in terms of disaster recovery performance and cost-efficiency. Prior studies have highlighted the benefits of cloud-based disaster recovery solutions, but often lack detailed platform-specific analyses. For instance, Abualkishik et al. [5] provided a general overview of cloud disaster recovery strategies without examining AWS and Azure performance metrics. Arogundade [6] compared cloud versus traditional disaster recovery methods but did not analyze AWS and Azure's disaster recovery services in depth. Alhazmi [7] proposed an adaptive disaster recovery model for SMEs but did not evaluate large-scale RTO or RPO outcomes.

This research aims to bridge these gaps by conducting a comparative performance, cost, and risk analysis of cloud-based disaster recovery solutions-specifically evaluating AWS and Microsoft Azure for Small and Medium Enterprises (SMEs). The study will assess key performance indicators-Recovery Point Objective (RPO), Recovery Time Objective (RTO), and Total Cost of Ownership (TCO)-under three simulated disaster scenarios: single-VM crashes, full-region outages, and network disruptions [8]. Each scenario is repeatedly executed on both AWS Elastic Disaster Recovery (EDR) and Azure Site Recovery (ASR) to gather robust comparative data. The analysis will also include network-level Quality of Service (QoS) metrics (latency, throughput, and packet loss) during these events, as well as a detailed cost breakdown into storage, compute (replication) and bandwidth components to provide a holistic evaluation of each platform's DR solution.

To improve contextual relevance, this study also incorporates SME-specific requirements from the Indonesian market, including analysis of regional infrastructure, security controls, and sector-specific DR readiness. Additionally, this study will integrate Risk Assessment and Business Impact Analysis (BIA) as core elements of disaster recovery planning. A threat modeling framework is adopted to quantify and prioritize risks based on likelihood and impact, using weighted scoring aligned with real-world SME scenarios. Key elements such as threat analysis, attack vectors, and attack surfaces will be examined to enhance the comprehensiveness of the DR strategy [9]. Furthermore, SME case studies are used to evaluate trade-offs in DR platform selection, highlighting contrasting preferences such as cost-efficiency vs. ease-of-use. By combining technical performance metrics with cost-efficiency analysis and risk evaluation, this study provides actionable guidance tailored to the needs of Indonesian SMEs and supports data-driven decision-making in selecting the most suitable cloud-based disaster recovery solutions.

1.2 Statement of Problem

The growing dependence of Small and Medium Enterprises (SMEs) on cloud services necessitates robust disaster recovery (DR) solutions to ensure business continuity. Amazon Web Services (AWS) and Microsoft Azure offer comprehensive DR capabilities, but choosing the most suitable platform requires a multidimensional evaluation of performance metrics, cost-efficiency, and risk management. Few existing studies perform an integrated, comparative analysis of these aspects for cloud DR in SMEs with limited IT budgets. This research conducts

a data-driven comparative analysis of AWS and Azure DR solution for SMEs. Specifically, it addresses:

1. How can the DR performance of Amazon Web Services (AWS) and Microsoft Azure be evaluated using metrics such as Recovery Time Objective (RTO), Recovery Point Objective (RPO), and failover efficiency under the defined disaster scenarios?
2. What is the comparative cost-efficiency of Amazon Web Services (AWS) and Microsoft Azure DR implementations, considering component like storage, compute (replication), bandwidth, and overall Total Cost of Ownership (TCO)?
3. How can a comprehensive risk and operational analysis including Business Impact Analysis (BIA) and Quality of Service (QoS) evaluation be applied to assess vulnerabilities, recovery responsiveness, and potential financial losses in AWS and Azure DR platforms, to support optimal platform selection for SMEs?

1.3 Research Objectives

The primary objective of this research is to provide a comprehensive, data-driven comparison of Amazon Web Services (AWS) and Microsoft Azure disaster recovery solutions in terms of performance, cost-efficiency, and risk management. The study assists Small and Medium Enterprises (SMEs) in selecting the most suitable cloud DR platform for business continuity. The specific objectives are:

1. To compare Amazon Web Service (AWS) and Microsoft Azure DR performance using metrics such as Recovery Time Objective (RTO) and Recovery Point Objective (RPO) across the defined failure scenarios, determining how each platform meets continuity requirements.
2. To evaluate the cost-efficiency of Amazon Web Service (AWS) and Microsoft Azure DR implementations by analyzing components costs (storage, compute, bandwidth) and total cost of ownership (TCO), identifying which platform offers lower operational expenses.
3. To synthesize the performance, cost, and risk findings including Quality of Service (QoS) metrics and Business Impact Analysis (BIA) into actionable recommendations that help SMEs choose the optimal cloud DR solution based on performance targets, budget constraints, and real-world operational impact.

1.4 Scope of the Research

The assumptions and limitations of this study are:

1. The research focuses on the implementation, performance evaluation, and risk assessment of DR solutions on Amazon Web Services (AWS) and Microsoft Azure using key metrics such as Recovery Time Objective (RTO), Recovery Point Objective (RPO), Total Cost of Ownership (TCO), and risk factors like threat analysis and attack vectors.
2. Cost-efficiency is analyzed via actual service pricing, but broader financial modeling (e.g., dynamic pricing tiers) is excluded.
3. The cloud model used in this study is Infrastructure as a Service (IaaS), utilizing virtual machines, cloud storage, and native DR services (AWS Elastic Disaster Recovery and Azure Site Recovery) for data replication and failover.
4. The DR testing environment simulates business-critical applications (databases, web servers) to evaluate continuity under disasters. Each defined disaster scenario (VM crash, region outage, network disruption) is repeatedly executed on both platforms to ensure consistency and reliability of results.
5. Performance metrics include RTO (recovery time), RPO (data loss tolerance), availability, and cost-efficiency (including storage, compute, bandwidth costs).
6. The cloud providers examined are Amazon Web Services (AWS) and Microsoft Azure. These two platforms were chosen because they are the dominant public cloud providers with mature, widely-used DR services, especially in the SME market segment. This focus allows for a deep, controlled comparison of two leading solutions. Other providers (e.g. Google Cloud Platform, IBM Cloud) are excluded from this study to limit scope and complexity; evaluating additional vendors would substantially expand the experiment without changing the core insights for SMEs.
7. The operating system for test is Linux Ubuntu 20.04 LTS with Apache web server and MySQL databases.
8. The research includes SME-focused risk scoring and Business Impact Analysis (BIA) based on simulated failure scenarios and recovery outcomes. Detailed threat modeling, attack vector mapping, and security auditing are out of scope.
9. On-premises infrastructure and physical hardware are not tested; the study focuses on cloud-only DR.
10. Monitoring tools (AWS CloudWatch, Azure Monitor, cost dashboards) are used for collecting performance and usage data.

1.5 Hypothesis

The hypothesis of this research is that implementing Disaster Recovery solutions on Amazon Web Services (AWS) and Microsoft Azure will yield robust recovery capabilities; however AWS is expected to achieve faster recovery (shorter RTO and RPO) due to its optimized replication and global infrastructure. AWS is expected to achieve shorter RTO and RPO (on the order of 10-15 minutes), whereas Azure may experience longer recovery times and higher operational costs. In terms of cost, AWS's pricing model is expected to result in lower overall DR expenses compared to Azure's, which tends to have higher charges for data replication and transfer. Both platforms are presumed to provide strong risk mitigation features, though AWS's broad global infrastructure may reduce downtime under severe conditions. Thus, AWS is hypothesized to be the more efficient choice for SMEs requiring rapid recovery and cost-effective DR solutions. The analysis is conducted impartially and with full academic rigor; no vendor is promoted. All findings are presented objectively, without favoring any provider.

1.6 Research Methodology

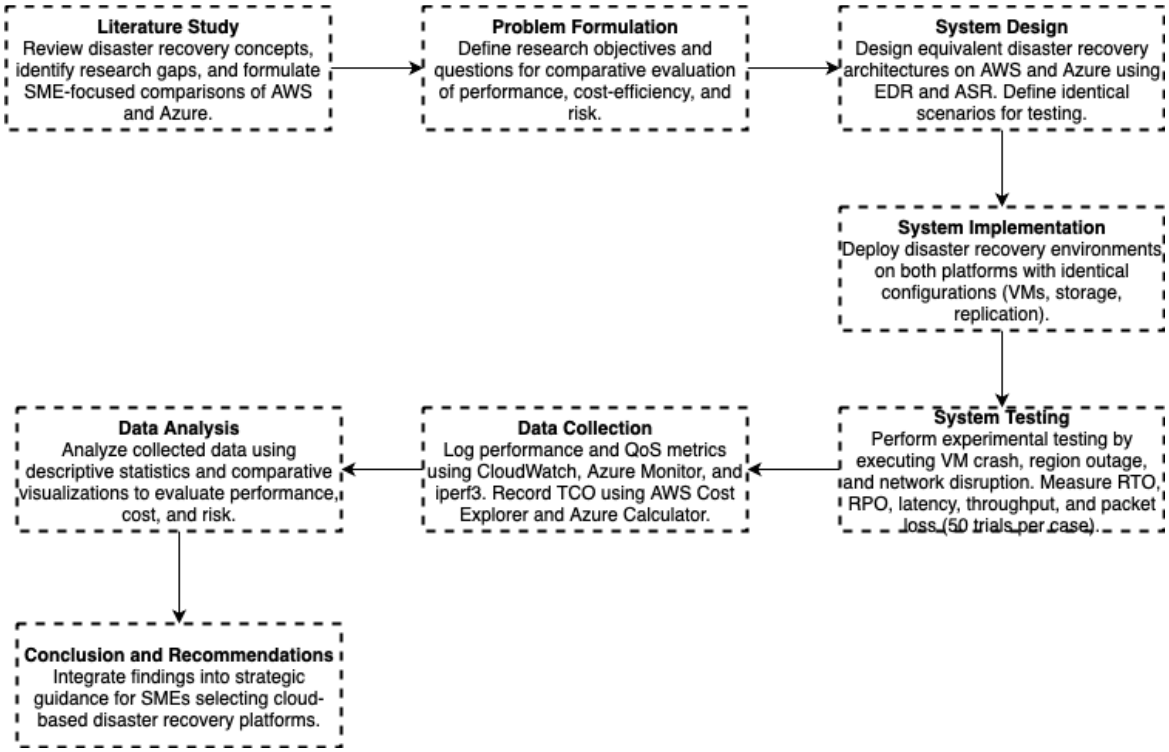

Figure 1. 1 Schema Methodology

This study adopts an experimental methodology to evaluate and compare the disaster recovery (DR) capabilities offered by Amazon Web Services (AWS) and Microsoft Azure,

with a focus on performance, cost, and risk analysis for Small and Medium Enterprises (SMEs). The research is structured into eight systematic stages as follows:

A. Literature Study

A comprehensive literature review was conducted to explore core DR concepts, with particular attention to Recovery Time Objective (RTO), Recovery Point Objective (RPO), Total Cost of Ownership (TCO), and Business Impact Analysis (BIA). The review identified knowledge gaps in empirical, side-by-side evaluations of AWS and Azure DR platforms, particularly those tailored for SMEs with limited IT budgets.

B. Problem Formulation

From the literature findings, the research problem was formulated to address the need for a comparative assessment of AWS and Azure disaster recovery solutions across three dimensions: performance, cost-efficiency, and risk mitigation. Specific research questions guide the evaluation of platform responsiveness under simulated disaster events, operational expenditure, and risk exposure.

C. System Design

Identical system architectures were established on both AWS and Azure to ensure a fair comparison. Each deployment used equivalent virtual resources (the same number of vCPUs, memory size, and disk capacity) and ran Ubuntu 20.04 LTS virtual machines with Apache web servers and MySQL databases. Both the AWS and Azure environments employed their native DR services (AWS Elastic Disaster Recovery and Azure Site Recovery). All configuration parameters for example, instance sizes, storage performance tiers, replication schedules, and network settings were matched across both platforms. This parity in setup (identical resource specifications and recovery objectives) guaranteed that both AWS and Azure environments were tuned for equivalent performance and reliability. The test scenarios included a VM crash, a full-region outage, and a network disruption.

Furthermore, additional validation was conducted on the Azure Site Recovery setup to ensure no critical disaster recovery features were unintentionally disabled. Replication frequency, storage redundancy, failover policy, and recovery plans were carefully reviewed and enabled to align with AWS Elastic Disaster Recovery's default capabilities. This ensured that Azure operated in a fully optimized state for disaster recovery and that any observed performance differences were the result of architectural behavior rather than configuration limitations.

To support this equivalency claim, a detailed comparison of system parameters was conducted and is presented in Table 1.1 below. The table outlines the specific resource and configuration attributes for both cloud platforms, along with validation notes that confirm each setting was aligned to ensure fairness in performance and cost evaluation.

Table 1. 1 System Parameter Validation between AWS and Azure on Identical Architecture

Parameter	AWS	Azure	Validation Notes
vCPU	4 vCPU	4 vCPU	Identical number of virtual CPUs
Memory	8 GB	8 GB	Identical RAM capacity
Storage	100 GB SSD (EBS gp2)	100 GB SSD (Premium SSD)	Identical SSD type and capacity
Operating System	Ubuntu Server 20.04 LTS	Ubuntu Server 20.04 LTS	Identical OS version
DR Solution	AWS Elastic Disaster Recovery (AWS Elastic DR)	Azure Site Recovery	Two equivalent cloud-native DR solutions
Replication Schedule	Continuous replication every 5 minutes	Continuous replication every 5 minutes	Identical data replication intervals
Networking	VPC with centralized subnet and security groups (CIDR /16)	Virtual Network (VNet) with subnet and NSG (CIDR /16)	Identical virtual network architecture (CIDR, subnets, security rules)

D. System Implementation

The DR environments were implemented using real cloud infrastructure. Replication, monitoring, and failover mechanisms were activated and configured with consistent parameters across both platforms. Services were deployed in the respective primary regions and replicated to geographically distant secondary regions to simulate cross-regional resilience.

E. System Testing

Each failure scenario was repeatedly tested on both platforms.. During each trial, failover and failback processes were executed to evaluate RTO and RPO. Additional performance metrics latency, throughput, and packet loss were recorded using tools like iperf3, AWS CloudWatch, and Azure Monitor.

F. Data Collection

Performance data were collected automatically through logs and dashboards. Cost data were obtained from AWS Cost Explorer and Azure Pricing Calculator, including charges for compute, storage, bandwidth, and licensing. Risk-related data were derived from failure impacts and modeled using Business Impact Analysis (BIA).

G. Data Analysis

The collected data were analyzed using descriptive statistics and comparative visualization techniques (boxplots, bar charts). RTO and RPO values were evaluated per scenario. TCO was broken down into key cost components. QoS indicators were used to assess network resilience, while risk analysis matrices ranked disaster scenarios by severity and probability.

H. Conclusions and Recommendations

Findings from the experimental testing and analysis stages were synthesized into strategic insights to support SME decision-making. The conclusion integrates both technical and financial perspectives, offering a comprehensive understanding of how cloud-based disaster recovery platforms perform under real-world failure scenarios. The recommendations provide practical guidance for selecting DR platforms based on key considerations such as performance targets, budget constraints, and risk tolerance thresholds relevant to business continuity planning.