International Journal of Computer Trends and Technology
ISSN: 2231–2803 / https://doi.org/10.14445/22312803/IJCTT-V72I2P108

Volume 72 Issue 2, 41-45, February 2024

Original Article

Implementing Infrastructure-as-Code with Cloud Disaster Recovery Strategies

Jairenz T. Batu

College of Computer Studies, Tarlac State University, Tarlac, Philippines

Corresponding Author : j.batu01233@student.tsu.edu.ph

Received: 04 January 2024 Revised: 05 February 2024 Accepted: 16 February 2024 Published: 29 February 2024

Abstract *- Data is a crucial asset for organizations in the modern digital age. Successful companies rely on information at every stage of their decision-making process. However, with the increasing importance of data comes a rise in threats to its security and integrity. These threats, such as software or hardware failures, natural disasters, or human errors, can be unplanned. To address these risks, regular backups are essential. While data recovery systems may function as intended, there is another challenge that often goes overlooked: determining where to restore the data. In this study, the researcher proposes combining a cloud data backup strategy like Warm Standby, Backup & Restore, and Pilot Light with infrastructure-as-code (IaC) implementation. This approach aims to reduce the risk of data loss and enable businesses to recover easily in case their main systems go offline due to ransomware attacks. IaC automates the provisioning and managing of infrastructure resources such as servers and networks through code. By leveraging IaC alongside cloud-based backup strategies, organizations can enhance their ability to protect critical data and ensure business continuity in challenging scenarios.*

Keywords *- Backup strategies, Business continuity, Cloud computing, Disaster recovery, Infrastructure-as-Code.*

1. Introduction

Data is critical to all organizations in the modern business landscape. [1] With the maturity of cloud computing technology, various organizations, enterprises, and government agencies are choosing databases as the main management systems for storing data. If the database fails, the effects on the organization are incalculable. Such failures also result in the complete shutdown of the organization's systems and services. [2]

In some circumstances, even with adequate safeguards, essential records can still be destroyed due to the possibility of human errors, hardware failure, or software failure. [1] In addition to these, more sinister forms of a disaster are those deliberately planned by humans either through trashing the data, stealing encryption keys, or through malicious attacks. [3]

One form of a malicious attack whose objective is to block its victims from accessing their valuable data and resources by locking or encrypting data stored on the target computer, making the data totally unusable, has captured the attention of cybersecurity experts in recent years due to the rising number of attacks. [4] These kinds of attacks are called ransomware. Ransomware is a form of extortion-based malware threat causing billions of dollars in losses for organizations. A ransom is demanded by the attackers from the victims before their data is released. [5]

One of the ways to prevent catastrophic situations like these is to execute regular backup operations. [1] Backup strategies are usually measured using Recovery Point Objective (RPO) and Recovery Time Objective (RTO) metrics. [6] A backup strategy that utilizes a sufficient time-delineated and low RPO capability is preferred to address ransomware attacks. [5]

Despite the advantages of the technology, such as the development of highly available storage systems and availability of cloud technologies, which have dramatically reduced recovery times, [7] the disaster recovery techniques still lack a technique to recover complex information systems if a catastrophe strikes. Assuming data recovery efforts begin, another frequently overlooked issue is where to restore the data. The dilemma is whether it is possible to restore the data to a secondary virtual infrastructure and fail over the traffic to it or if a completely new infrastructure is required, depending on the impact and severity of the disaster.

The cloud enables organizations to design flexible and scalable backup recovery plans, allowing them to meet their

RTO and RPO targets. [5] In the cloud, users can automate the provisioning of cloud resources, like servers and networks, using a technique called Infrastructure-as-Code (IaC). [8] Mixing the cloud with Infrastructure as Code (IaC) techniques permits the organization to restore not only its data backups but also its underlying virtual machines, database systems, firewalls, and network subnets by defining and representing the infrastructural entities through codes driving the restoration procedure to be fully automated. [9]

Using this method, the previously designed and deployed infrastructure using a graphical user interface can now be versioned and maintained in coded format while also serving as a comprehensive backup of the infrastructure architecture. [10] Backup processes can be automated with infrastructure as code, assuring the continuity and availability of enterprise services and activities. [11] Designing a fully automated backup and disaster recovery for an information system and its underlying infrastructure allows organizations to achieve a low RTO and RPO goal, allowing them to combat various threats and catastrophes.

2. Methodology

In this research, different backup restoration and recovery strategies and procedures will be simulated by the proponent to test the amount of time needed to fully recover a web application. The backup strategies utilized are Backup and Restore, Pilot Light, and Warm Standby. In the backup and restore strategy, the backups are created on an external storage service and restored when needed. Unlike other strategies, no secondary system is available on standby and will only be started during the restoration process. Backup and restore are associated with higher RTO (recovery time objective) and RPO (recovery point objective). The term pilot light is often used to describe a disaster recovery scenario in which a minimal version of an environment is always running in the cloud. The idea of the pilot light is an analogy that comes from the gas heater. In a gas heater, a small flame that's always on can quickly ignite the entire furnace to heat a house. Finally, the term warm standby describes a disaster recovery scenario in which a scaled-down version of a fully functional system is always running in the cloud. [12]

Disaster recovery plans and backup strategies are usually measured using the Recovery Point Objective (RPO) and Recovery Time Objective (RTO) metrics. [6] RPO is the acceptable difference between two back-ups based on the time they were created, while RTO defines the maximum reasonable time a service or system may be offline. [7] The RPO and RTO combined answer the questions of how long the system should be online again to start serving requests and how the system should look once it's up. [5] Based on the study carried out by Thomas Galligher in 2018, the best method for achieving the level of stability required to mitigate issues caused by ransomware attacks is those backups done incrementally that would allow a system to go back to any given day to restore the system. To measure RTO, a stopwatch will be utilized to record the time from the point of the disaster until the point of recovery. The stopwatch will stop as soon as the web application loads in the browser. As for the RPO, a scheduler such as CRON will trigger the backup schedules.

An illustration of the RPO and RTO is shown in Figure 1.

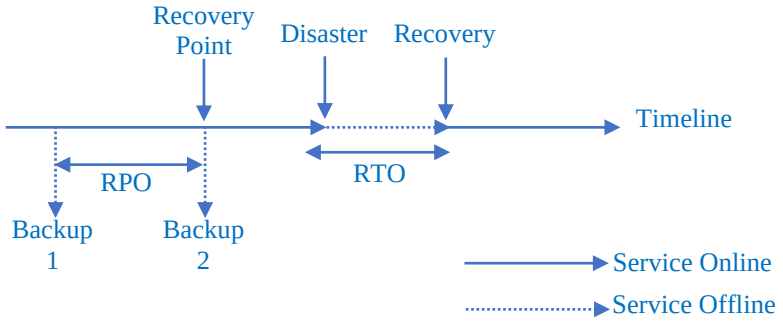

Fig. 1 RTO and RPO in the context of an IT Service

Together with the RPO and RTO data, the costs of each strategy will also be discussed in a matrix, as shown in Table 1. The goal of these tests is to prove the possibility of creating a secondary virtual infrastructure that will act as a fail-over system when an intruder attacks the primary infrastructure without the need to exert too much budget from the organization implementing the business continuity plan. These tests will be facilitated using a three-tier web application model.

Table 1. A Comparison of Backup Strategies

Backup Strategy	Backup & Restore	Pilot Light	Warm Standby
RPO (minutes)	00	00	00
RTO (minutes)	00	00	00
Implementation Cost	$$$	$$$	$$$

2.1. Backup Strategies with Manual Restoration Process

In the first part of this methodology, the backup strategies Backup and Restore, Pilot Light, and Warm Standby [12] will be combined into a manual restoration process. These backup procedures directly correlate to the RPO metrics. On the other hand, the restoration, which correlates to the RTO metric, will be done manually, including the restoration of the last known resource configuration of the virtual infrastructure.

2.2. Backup Strategies with Automated Restoration Process through IaC

In the second phase of the tests, Infrastructure as Code (IaC), which is an approach to automate the creation of infrastructure and related entities [13], will be integrated into the mix, which should supposedly lower the RTO metric. All three cloud backup strategies will be utilized, but the restoration process will be carried out using an IaC tool called Terraform. This tool is added to the mix to test if the research will achieve its aim of producing the lowest RPO and RTO for the system to decrease the business impact caused by downtime.

2.3. Cloud Architecture and Design

In this study, the system restoration scenario is envisaged to occur within a different architecture with independent subnets and network configurations. The goal is to avoid cross-contamination from other systems infected by the same ransomware. Figure 2 depicts the chosen cloud environment. The deliberate isolation extends to unique subnets and network configurations to improve the recovered system's resistance against lingering threats while leveraging the cloud infrastructure's design with best practices to provide accessibility, resource optimization, and speedy deployment. The strategy reduces the possibility of lateral movement inside the network and strengthens the system against future re-infection.

The cost of implementing each strategy will be determined mostly by the number of resources required to support each method. The researcher computed an estimate, through the help of the Amazon Web Services' pricing calculator, for each strategy based on AWS' advertised on-demand resource rates. The researcher's calculation for backup techniques requiring standby resources is based on a monthly billing rate of 750 hours.

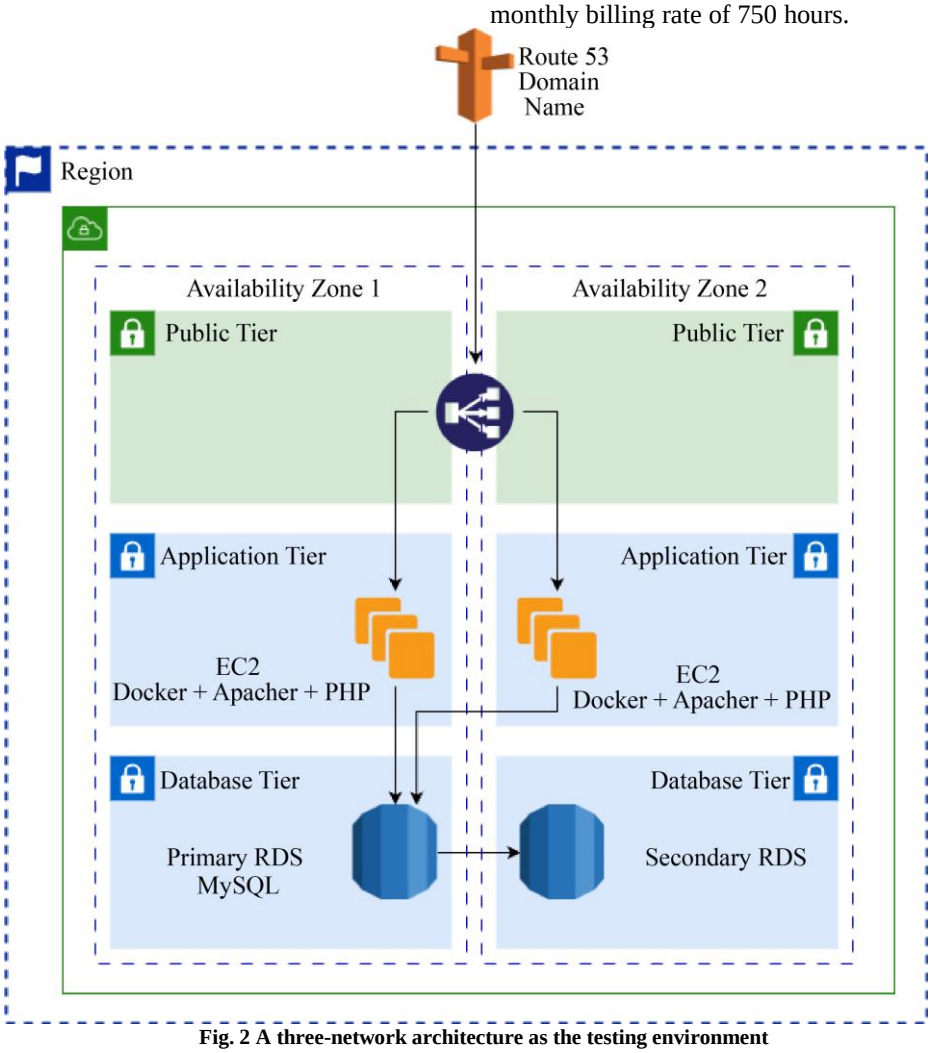

Fig. 2 A three-network architecture as the testing environment

This network design will be used in all the tests listed above, including both manual and automated testing methodologies. The entire network architecture must be restored to construct a new virtual private cloud, complete with a private network subnet, internet gateway, security groups, route tables, application load balancer, database profiles, and virtual machines. The table below shows which cloud resources are expected to be created or modified for each backup and restoration strategy.

Table 2. A list of cloud resources that will be deployed or modified upon restoration

Cloud Resources	Backup Strategies		
	Backup & Restore	Pilot Light	Warm Standby
VPC Subnets	✓		
Internet Gateway	✓		
Route Table	✓		
Security Groups	✓		
Database	✓		
Application Load Balancer	✓	✓	
Auto Scaling Group	✓	✓	
DNS record	✓	✓	✓

3. Results and Discussion

This research tested three different backup and recovery strategies, compared their RPO and RTO results, and analyzed the estimated cost to implement each strategy. The first part utilized a manual process for the restoration, while the second phase included infrastructure as code to lower the RTO targets. Details of the results can be seen in the next part.

3.1. Backup Strategies with Manual Restoration Process

In the first phase of the tests, the results show that the Warm Standby and Pilot Light strategies demonstrated the lowest RTO and RPO. The RPO can be lowered to just five minutes before the backup procedures affect any production systems and other processes running on the server. The Backup and Restore procedure would be possible within a 1 hour RPO target. This strategy achieves the highest RPO target.

Meanwhile, the RTO for the pilot light strategy tends to be between 10 and 12 minutes, and this would mainly depend on the user's knowledge of executing the scaling procedures and adding new instances and resources to support production workloads. Warm standby's RTO can be as low as less than 2 minutes since instances are already in their running state, as expected based on Table 2, and the user would only need to update routing configurations (e.g., DNS updates). Finally, the Backup and Restore strategy was last on the list, as it took 53-60 minutes before system restoration was completed, as most or all cloud resources must be created manually by the administrator.

Table 3. A Comparison of backup strategies

Backup Strategy	Backup & Restore	Pilot Light	Warm Standby
RPO (minutes)	60	5	5
RTO (minutes)	53-60	10-12	<2
Implementation Cost	$$	$$$	$$$$$

In terms of cost, Warm Standby is the most expensive option because all of the required cloud resources have already been deployed and are simply waiting in the background for incoming traffic via DNS routing updates. Pilot Light is around 40% less expensive, while Backup and Restore are 60% less expensive.

3.2. Backup Strategies with Automated Restoration Process through IaC

For the second phase of the tests, an Infrastructure as Code tool called Terraform will be added to the mix to implement the automation of the recovery process, which should lower the RTO of the backup strategies. Terraform had the most effect on the Backup and Restore strategy. The RTO was lowered to just 8-10 minutes, a significant 85% faster than the manual process. It had the least effect on the Warm Standby strategy, as it no longer needed to do much since most resources already existed in this setup. What's interesting is the effect on the Pilot Light strategy. Due to the advantages of automated processes, the RTO has been lowered to 4-6 minutes.

The backup strategies tested in both Phases 1 and 2 had the same cost and budget needs because the strategy did not change, and the IaC employed was open source, which does not require a license or subscription.

Table 4. A Comparison of Backup Strategies with IaC Implementation

Backup Strategy	Backup & Restore	Pilot Light	Warm Standby
RPO (minutes)	60	5	5
RTO (minutes)	8-10	4-6	<2
Implementation Cost	$$	$$$	$$$$$

4. Conclusion

To conclude this study, it was understood that the RPO metric had an indirect relationship with the cost needed to implement the strategy. A lower RPO target (minutes compared to hours) means a higher cost for the business. If a business wants a thorough disaster recovery and business continuity plan, it must implement a strategy with the lowest RPO and thus boost its budget accordingly. They should also explore controlling their cloud infrastructure through IaC, which would provide them with the option of using the Warm Standby Strategy or the Pilot Light. Meanwhile, for the RTO metric, it was proven that integrating IaC with the backup strategies would dramatically lower the RTO target without increasing the cost. Moreover, based on the data gathered, it can be concluded that the warm standby strategy is the most applicable strategy for ransomware attacks; however, due to cost, the pilot light strategy is a good alternative. These results would be beneficial for organizations and different companies that would like to balance all three metrics and protect their organizational data from disasters and catastrophes.

References

[1] Qun Li, and Honglin Xu, "Research on the Backup Mechanism of Oracle Database," *2009 International Conference on Environmental Science and Information Application Technology*, Wuhan, China, pp. 423-426, 2009. [CrossRef] [Google Scholar] [Publisher Link]

[2] Dan Jin, and Qiong Wang, "CDP Backup and Recovery Method for Ensuring Database Consistency," *2021 IEEE International Conference on Power Electronics, Computer Applications (ICPECA)*, Shenyang, China, pp. 722-728, 2021. [CrossRef] [Google Scholar] [Publisher Link]

[3] Tejinder Pal Singh Brar, Dhiraj Sharma, and Sawtantar Singh Khurani, "Disaster Recovery and Business Continuity Planning for Electronic Banking: A Comparative Study," *International Journal of Commerce and Management*, vol. 9, pp. 54-71, 2015. [Google Scholar] [Publisher Link]

[4] SH Kok et al., "Ransomware, Threat and Detection Techniques: A Review," *International Journal of Computer Science and Network Security*, vol. 19, no. 2, pp. 136-146, 2019. [Google Scholar] [Publisher Link]

[5] Jason Thomas, and Gordon Galligher, "Improving Backup System Evaluations in Information Security Risk Assessments to Combat Ransomware," *Computer and Information Science*, vol. 11, no. 1, pp. 14-25, 2018. [CrossRef] [Google Scholar] [Publisher Link]

[6] Mugoh Leon, Ateya Ismail Lukandu, and Shibwabo Bernard Kasamani, "Continuous Data Protection as a Strategy for Reduced Data Recovery Time," *Journal of Systems Integration*, vol. 2, pp. 54-69, 2011. [Google Scholar]

[7] Alexander Lenk, and Stefan Tai, "Cloud Standby: Disaster Recovery of Distributed Systems in the Cloud," *European Conference on Service-Oriented and Cloud Computing*, Berlin, Heidelberg, vol. 8745, pp. 32-46, 2014. [CrossRef] [Google Scholar] [Publisher Link]

[8] Tiago Espinha Gasiba et al., "Raising Security Awareness of Cloud Deployments Using Infrastructure as Code through CyberSecurity Challenges," *ARES '21: Proceedings of the 16th International Conference on Availability, Reliability and Security*, pp. 1-8, 2021. [CrossRef] [Google Scholar] [Publisher Link]

[9] Orest Lavriv et al., "Method of Cloud System Disaster Recovery Based on "Infrastructure as a Code" Concept," *2018 14th International Conference on Advanced Trends in Radioelecrtronics, Telecommunications and Computer Engineering (TCSET)*, Lviv-Slavske, Ukraine, pp. 1139-1142, 2018. [CrossRef] [Google Scholar] [Publisher Link]

[10] Manish Kumar et al., "Infrastructure as Code (IaC): Insights on Various Platforms," *Sentiment Analysis and Deep Learning*, vol. 1432, pp. 439-449, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[11] Sneh Pandya, and Riya Guha Thakurta, "Business Solutions with Infrastructure as Code," *Introduction to Infrastructure as Code*, Berkeley, CA: Apress, pp. 83-96, 2022. [CrossRef] [Google Scholar] [Publisher Link]

[12] Glen Robinson, Attila Narin, and Chris Elleman, "Using Amazon Web Services for Disaster Recovery," *Amazon Web Services-Using AWS for Disaster Recovery*, pp. 1-22, 2014. [Google Scholar] [Publisher Link]

[13] Kief Morris, *Infrastructure as Code: Dynamic Systems for the Cloud Age*, O'Reilly, pp. 1-399, 2021. [Google Scholar] [Publisher Link]