MODELING AND ASSESSMENT OF RELIABLE CLOUD STORAGE SOLUTIONS FOR CULTURAL HERITAGE ASSETS

Saso Nikolovski
AUE University, Faculty of Informatics-Skopje, North Macedonia, sasho.nikolovski@fon.edu.mk
Anita Ilieva Nikolovska
Macedonian Academy of Sciences and Arts, North Macedonia, anita@manu.edu.mk
Bozidar Milenkovski
Faculty of Information and Communication Technologies-Bitola, North Macedonia, bozidar.milenkovski@uklo.edu.mk

Abstract: This paper presents a comprehensive framework for evaluating the performance and reliability of a fully cloud-hosted disaster-recovery solution dedicated to the long-term preservation of cultural heritage data. The study focuses on Microsoft Azure Recovery Service-MARS, a platform that provides continuous, encrypted backup and granular recovery options designed to protect critical information against natural disasters, cyberattacks, and hardware failures. To ensure realistic results, we conducted year-long empirical measurements in a live production environment that mirrors the operational conditions of institutions responsible for safeguarding digital cultural assets. Backup and restore jobs were executed on a dedicated virtual machine connected through a symmetrical 200/200 Mbps Internet link, allowing the collection of detailed metrics such as average data throughput, Recovery Point Objective-RPO, Recovery Time Objective-RTO, and monthly service cost. These measurements capture the true performance envelope of the cloud service, including the impact of network fluctuations and encryption overhead.

Alongside performance monitoring, a system-dynamics reliability model was developed to quantify end-to-end service resilience.

The model represents the on-premises data center, the Internet service provider-ISP, and the Azure cloud as a series configuration and applies classical reliability theory to calculate time-dependent availability. Mean time between failures and published service level agreements-SLA were used to parameterize the simulation.

Results reveal that while Azure infrastructure consistently delivers near-perfect service availability (approaching eleven nines of uptime) the overall system reliability is ultimately governed by the weakest component of the delivery chain.

Over a 15-day observation window the integrated system achieved an aggregate reliability of approximately 0.97, with the ISP connection identified as the dominant limiting factor. These findings provide a practical foundation for designing resilient digital repositories for cultural heritage institutions. They emphasize the need for redundant high-bandwidth connectivity and continuous validation of recovery procedures to maintain strict RPO and RTO objectives. When such network safeguards are combined with the security, encryption, and long-term retention features of Microsoft Azure Recovery Service, organizations can confidently implement a sustainable, standards-compliant archival strategy capable of protecting irreplaceable cultural and historical artifacts for decades to come.

Keywords: Reliability, MARS, Azure, disaster recovery, cultural heritage

1. INTRODUCTION

Safeguarding digital representations of cultural and historical artifacts requires protection strategies that remain effective for decades. Absolute zero-downtime is unattainable because of hazards such as natural disasters and cyberattacks, so strong Business Continuity Plans-BCPs and Disaster Recovery Plans-DRPs are essential (Mathew & Mai, 2018; Tamimi et al., 2019). Cloud technologies have become central to these efforts, and numerous studies have examined their economic feasibility, privacy safeguards, and reliability (Wood et al., 2010; Rebah & Sta, 2016; Mendonca et al., 2019; Mendonca et al., 2020). Yet many rely on laboratory simulations, providing only partial validation (Nikolovski & Mitrevski, 2023; Nikolovski et al., 2024).

To close this gap, we performed year-long measurements in a production data center where MARS was deployed. Performance metrics such as backup/restore throughput, RPO, and RTO were collected directly from system agents. In parallel, a system-dynamics model evaluated the reliability of the on-premises data center, ISP, and Azure cloud treated as a series configuration (Joshi & Sivalingam, 2014; Miller & Engemann, 2014; Tian et al., 2020; Liu et al., 2022; Rajini et al., 2020). This combined perspective provides a realistic foundation for selecting cloud storage strategies suited to long-term cultural heritage preservation.

2. RELATED WORK

Research on DRaaS highlights cloud infrastructure as a dependable and cost-effective option for critical data protection (Mathew & Mai, 2018). Comparative studies reveal the economic and operational advantages of cloud-centric recovery models while noting risks such as vendor lock-in and service outages (Wood et al., 2010; Rebah & Sta, 2016). Analytical and stochastic models further quantify recovery metrics like RPO and RTO, though many were validated only in laboratory settings (Mendonca et al., 2019; Mendonca et al., 2020).

Recent work stresses the need for live-data-center testing to obtain empirical values for key parameters (Nikolovski & Mitrevski, 2023). Microsoft Azure Recovery Service (MARS) fits this trajectory by offering near-continuous availability and granular recovery options (Liu et al., 2022; Rajini et al., 2020). Studies combining reliability theory with real performance data use system-dynamics and series-configuration models to show how failures in network links or cloud services influence end-to-end reliability (Nikolovski et al., 2024).

3. METHODOLOGY

Production environment - The study was carried out in a high-availability data center (Fig. 1) with three server nodes in a failover cluster connected to a shared Cluster Shared Volume-CSV. All protected workloads, including the measurement virtual machine-VM, operate as virtual instances. Network connectivity to Azure is via a symmetrical 200/200 Mbps Internet link, reflecting enterprise-class constraints (Joshi & Sivalingam, 2014; Nikolovski et al., 2025).

Fig. 1 High-availability data center

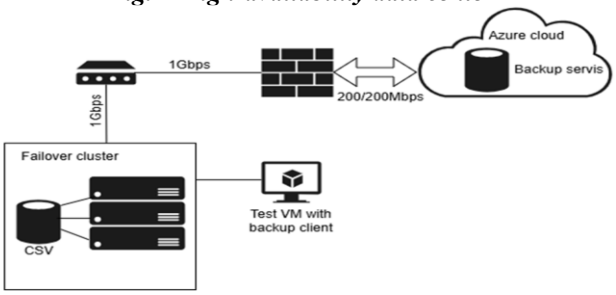

Source: The Author

Data protection setup - Within this environment, the MARS agent was deployed on the test VM to perform automated backup and restore operations (Fig. 2). The agent captures both the full system state and file-level data, enabling granular recovery of individual files or complete system images as required. Backup jobs execute in two phases: the first secures the operating-system and configuration state, while the second transfers application and file data to Azure storage. All backup data is encrypted in transit and at rest, consistent with the security standards of the service.

Measurement approach - A reverse-engineering technique was applied to extract performance metrics directly from the MARS agent logs during one year of scheduled operations.

Fig. 2 System dynamics model of MARS concept in production

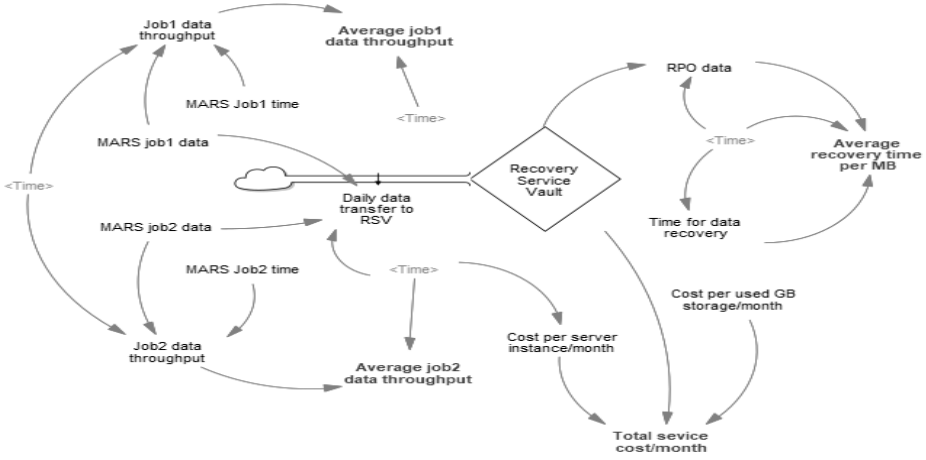

Source: The Author

From these logs we derived average data throughput (MB/s) and average recovery time per megabyte of restored data, as well as the total monthly service cost associated with the protected workload. These empirically determined variables serve as inputs for the subsequent modeling of reliability and system dynamics.

Reliability modeling - The cloud-only architecture (

Fig. ***3***) was modeled as a series of three components: on-premises data center, ISP link, and Azure cloud service (Miller & Engemann, 2014; Tian et al., 2020).

Fig. 3 System components for reliability modeling

Source: The Author

Overall system reliability R_{sys} (1) is the product of the component reliabilities, where each R_i is a time-dependent function of failure rate λ and mean time between failures-MTBF. Operational periods were set to 61,320 h (≈7 years) for data center and cloud, and 17,520 h (≈24 months) for the ISP.

$$R_{sys} = \prod_{i=1}^{n} R_i = \min_{1<i<n} R_i \quad (1)$$

A Service Level Agreement-SLA of 99.95 % was adopted for the ISP component for accessing cloud services in Azure. Using these parameters, a system-dynamics model with causal loop diagrams (

Fig. ***4***) captured the interactions among variables such as component failure, downtime, and SLA-driven availability. Simulation outputs highlights how a single weak link (in this case the ISP) can significantly reduce end-to-end reliability even when the cloud service itself offers near-perfect availability.

Fig. 4 Reliability model with causal loop diagram

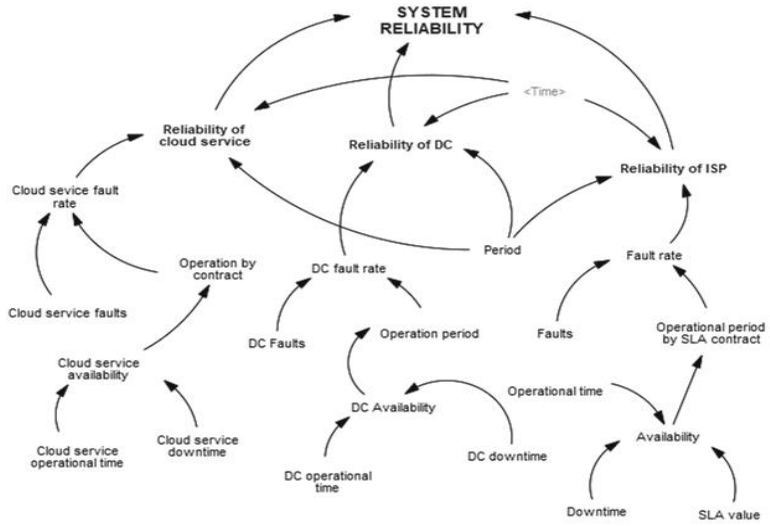

Source: The Author

4. RESULTS AND DISCUSSION

Performance outcomes - One year of continuous MARS operation showed that backup throughput was limited by the 200 Mbps link, with actual rates consistently below the theoretical maximum because of packet delays and competing traffic. The two-phase backup process confirmed that the second phase (user-data transfer) dominated total duration. Recovery operations displayed similar bandwidth constraints, and the mean recovery time per megabyte showed that large-scale restores are primarily network-bound despite Azure's advertised eleven-nines availability.

MARS simulation analysis –

Table ***I*** presents both the derived variables from the model and the simulation results obtained with the MARS configuration. The model (Fig. 2) shows an average Job 1 data throughput of 2.58 MB/s and Job 2 throughput of 1.83 MB/s, illustrating the difference between the initial system-state backup and the larger user-data transfer phase. The mean recovery time of 5.57 s/MB confirms that large-scale restores are primarily governed by network latency and the overhead of encryption and data rehydration.

Table I Results from the MARS concept simulation

Variable	Value
Derived variables from the model	
Average job1 data throughput (MB/s)	2.57731
Average job2 data throughput (MB/s)	1.83045
Average recovery time per MB (sec)	5.57246
Model simulation with test data	
Test data (MB)	531,012
Backup time Job1(Test data) (hours)	57.2315
Backup time Job2(Test data) (hours)	80.5831
Recovery time (Test data) (hours)	26.47

Source: The Author

The extended simulation with a test dataset of 531 GB highlights the cumulative impact of these rates. The first backup phase (Job 1) required approximately 57.2 hours, while the second phase (Job 2) consumed an additional 80.6 hours, reinforcing the earlier observation that the user-data transfer dominates total backup duration. The full recovery time of about 26.5 hours is consistent with the empirically measured values and supports the conclusion that the available 200 Mbps link, rather than Azure's internal infrastructure, is the primary performance constraint.

These results align with the reliability modeling discussed earlier, where the Internet service provider was identified as the critical bottleneck limiting overall system reliability to roughly 0.97 over the 15-day observation window. In practical terms, the table demonstrates that while the MARS platform provides near-continuous availability, network capacity planning and redundant high-bandwidth connections remain essential for institutions charged with safeguarding cultural heritage assets and meeting strict RPO/RTO objectives.

Reliability simulation - Using empirically derived parameters, the system (data center, ISP, and Azure cloud) was evaluated as a series configuration. Calculated component reliabilities were: Azure cloud 0.993952, data center 0.993952, and ISP 0.978981, yielding an overall reliability of approximately 0.967 for a 15-day window, calculated from expression (2). The ISP link proved to be the critical weak point even when the cloud service itself maintained near-perfect availability.

$$R_{sys} = R_{Azure} * R_{DC} * R_{ISP} \quad (2)$$

Implications for cultural heritage preservation - For organizations responsible for safeguarding cultural heritage assets (where data integrity and very long retention periods are critical) these findings underscore the need to evaluate the "weakest link" in the service chain. Even when a cloud provider offers near-perfect uptime, network interruptions or service-level breaches by the ISP can compromise recovery objectives. The results suggest that when RPO and RTO requirements are stringent, additional measures such as redundant high-bandwidth connections or geographically diverse network providers are essential to meet business continuity targets. Overall, the empirical measurements and reliability modeling confirm that Microsoft Azure Recovery Service is capable of delivering highly dependable long-term data protection. However, achieving end-to-end resilience for cultural heritage applications requires holistic planning that includes the supporting network infrastructure as an integral component of the disaster-recovery architecture.

Application to cultural heritage - Digitized cultural heritage collections (high-resolution images, 3-D scans, and extensive metadata) demand storage systems that guarantee integrity over decades while supporting timely restoration. This study demonstrates that a fully cloud-hosted solution such as MARS satisfies these requirements when paired with carefully planned connectivity and service-level agreements. Because heritage data often involves large files and specialized formats, the observed bandwidth limits directly affect RPO and RTO. Azure's near-perfect availability, encryption, and granular recovery features meet the long-term archival needs of museums, libraries, and archives.

5. CONCLUSION

This study presents an integrated evaluation of a disaster-recovery concept entirely deployed in the cloud, using MARS as the reference platform. By combining one year of real operational measurements with a system-dynamics

reliability model, we obtained a realistic view of the service's ability to protect mission-critical data (in this case cultural heritage data) and to meet long-term preservation requirements for cultural heritage assets. Empirical results confirm that Azure's infrastructure delivers extremely high availability yet the overall end-to-end reliability of the solution depends strongly on external factors, particularly the stability of the Internet service provider. The calculated system reliability of approximately 0.97 for a 15-day observation window illustrates how a single weaker component can limit the effectiveness of the entire disaster-recovery chain. For safeguarding cultural heritage, these findings highlight two key recommendations. First, network connectivity must be treated as a critical element of the architecture, warranting redundant high-bandwidth links or geographically diverse providers to satisfy strict RPO and RTO targets. Second, while Azure's encryption, granular recovery capabilities, and long-term data retention features meet stringent preservation standards, the broader disaster-recovery plan should integrate continuous monitoring and periodic validation of recovery procedures to maintain compliance and readiness over decades of operation. In conclusion, a fully cloud-based disaster-recovery solution built on Microsoft Azure Recovery Service provides a robust foundation for the protection and long-term preservation of digital cultural heritage. When paired with resilient network infrastructure and proactive management, such an approach offers a sustainable, high-reliability strategy for safeguarding irreplaceable cultural assets well into the future.

REFERENCES

Joshi, S. C., & Sivalingam, K. M. (2014). Fault tolerance mechanisms for virtual data center architectures. *Springer Science+Business Media*, New York.

Liu, Z., et al. (2022). Reliability modelling and optimization for microservice-based cloud application using multi-agent system. *IET Communications, 16*(10), 1182–1199.

Mathew, A., & Mai, C. (2018). Study of various data recovery and data backup techniques in cloud computing & their comparison. *Proceedings of the 3rd IEEE International Conference on Recent Trends in Electronics, Information & Communication Technology*, Bangalore, India.

Mendonca, J., Lima, R., Andrade, E., Araujo, J., & Kim, D. S. (2020). Multiple-criteria evaluation of disaster recovery strategies based on stochastic models. *16th International Conference on the Design of Reliable Communication Networks (DRCN)*, IEEE.

Mendonca, J., Lima, R., Queiroz, E., Andrade, E., & Kim, D. S. (2019). Evaluation of a backup-as-a-service environment for disaster recovery. *IEEE Symposium on Computers and Communications (ISCC)*, Barcelona, Spain.

Miller, H. E., & Engemann, K. J. (2014). Using reliability and simulation models in business continuity planning. *International Journal of Technology, Policy and Management, 5*(1), 43–56.

Nikolovski, S., Milenkovski, B., & Petreska, A. (2025). Reliability analysis through system dynamics modeling for intellectual capital protection with a cloud services. *60th International Scientific Conference on Information, Communication and Energy Systems and Technologies (ICEST)*, IEEE.

Nikolovski, S., Mitrevski, P., & Petreska, A. (2024). Performance analysis of cloud service-based data protection systems. *59th International Scientific Conference on Information, Communication and Energy Systems and Technologies (ICEST)*, IEEE.

Nikolovski, S., & Mitrevski, P. (2023). Data protection and recovery performance analysis of cloud-based recovery service. *58th International Scientific Conference on Information, Communication and Energy Systems and Technologies (ICEST)*, IEEE.

Rajini, N., et al. (2020). Reliability of cloud services provided to non-banking financial institutions. *International Journal of Control and Automation, 13*(2s), 165–172.

Rebah, H. B., & Sta, H. B. (2016). Disaster recovery as a service: A disaster recovery plan in the cloud for SMEs. *Global Summit on Computer & Information Technology*, Sousse, Tunisia.

Shahzadi, S., Ubakanma, G., Iqbal, M., & Dagiuklas, T. (2018). Autonomous, seamless and resilience carrier cloud brokerage solution for business contingencies during disaster recovery. *IEEE 20th International Conference on High Performance Computing and Communications*, Exeter, UK.

Tamimi, A. A., Dawood, R., & Sadaqa, L. (2019). Disaster recovery techniques in cloud computing. *IEEE Jordan International Joint Conference on Electrical Engineering and Information Technology (JEEIT)*, Amman, Jordan.

Tian, Y., Tian, J., & Li, N. (2020). Cloud reliability and efficiency improvement via failure risk based proactive actions. *Journal of Systems and Software, 163*, 110524.

Wood, T., Cecchet, E., Ramakrishnan, K., Shenoy, P., Van Der Merwe, J., & Venkataramani, A. (2010). Disaster recovery as a cloud service: Economic benefits & deployment challenges. *2nd USENIX Workshop on Hot Topics in Cloud Computing*, Boston, MA.