

Synthetic Users Disaster Recovery Plan

Document Version

- **Version:** 1.1
- **Date:** February 2026
- **Prepared by:** Artur Ventura, CTO & CISO

Plan Overview

- **Purpose:** To ensure rapid and efficient recovery of Synthetic Users' operations in the event of a disaster, specifically focusing on critical dependencies like Render and OpenAI.
- **Scope:** This DRP covers processes and protocols for switching operations from Render to AWS in the event of a Render failure and from OpenAI to Anthropic in the event of an OpenAI failure.

Critical Dependencies

- **Primary Dependencies:** Render for application hosting, AWS for core infrastructure and data storage, OpenAI for AI model integrations.
- **Secondary Options:** AWS (direct) for application hosting failover, Anthropic for AI model integrations.

Recovery Objectives

- **Recovery Time Objective (RTO):** The maximum acceptable time to restore critical functions after a disaster.
 - **Render to AWS Migration:** 4 hours
 - **OpenAI to Anthropic Switch:** 2 hours
- **Recovery Point Objective (RPO):** The maximum acceptable amount of data loss measured in time.
 - **Data Backup:** 1 hour
- **Maximum Tolerable Downtime (MTD):** The absolute maximum duration the business can sustain a disruption before unacceptable impact occurs.

- **SaaS Operations:** 8 hours
- **AI Model Integrations:** 6 hours

Disaster Recovery Protocols

Render Failure - Migration to AWS

1. **Detection and Assessment:** Monitor and quickly identify service disruption on Render. Confirm the outage's scope and expected duration.
2. **Activation of AWS Environment:**
 - Pre-configured AWS environments should be maintained, mirroring the Render setup.
 - Initiate the AWS environment, ensuring all services and databases are operational.
3. **Data Migration:**
 - Last data backup from AWS RDS or S3 (since data is already stored on AWS) is verified and confirmed accessible.
 - Ensure the RPO of 1 hour is met by verifying data integrity.
4. **DNS Update:**
 - Update DNS records to point to the AWS environment, minimizing the switch-over time to meet the RTO of 4 hours.
5. **Verification and Monitoring:**
 - Conduct thorough testing to confirm operational functionality on AWS.
 - Monitor performance and stability closely following the switch.

OpenAI Failure - Switch to Anthropic

1. **Detection and Assessment:** Identify failure in OpenAI services impacting operations. Assess the impact on service offerings.
2. **Switch to Anthropic:**
 - Pre-configure Anthropic models to match the functionality provided by OpenAI models closely.
 - Redirect API calls from OpenAI to Anthropic, ensuring minimal changes to the integration layer.
3. **Verification and Adjustment:**

- Test the integration thoroughly to ensure that Anthropic models perform as expected.
- Adjust configurations as needed to optimize performance and accuracy.

4. Communication:

- Inform internal teams about the switch to manage expectations and provide updated documentation if necessary.
- Notify key clients of the change, emphasizing the continuity of service and quality.

Post-Recovery Actions

- **Review and Analysis:** After the recovery, conduct a detailed review to analyze the response's effectiveness, documenting lessons learned.
- **Plan Update:** Update the DRP based on feedback and any changes in the technological landscape or business requirements.

Testing and Maintenance

- **Annual DRP Testing:** Simulate disaster scenarios annually to test the effectiveness of the DRP, focusing on the switch from Render to AWS and OpenAI to Anthropic.
- **DRP Updates:** Review and update the DRP semi-annually or following significant changes in technology or business operations.

Documentation and Training

- **DRP Document:** Maintain a comprehensive, accessible DRP document detailing all protocols, procedures, and recovery objectives.
- **Training:** Regularly train relevant staff on DRP protocols, ensuring clear understanding and readiness to act in the event of a disaster.