

Synthetic Users Business Continuity & Technology Resiliency Exercise 2026

Version: 1.0

Date of Exercise: March 26, 2026

Prepared by: Artur Ventura, CTO & CISO

Classification: Internal — Confidential

CRA Reference: 15.2.4, 15.2.5, 15.3.2, 17.4.2

1. Exercise Overview

1.1 Objective

To validate the effectiveness of Synthetic Users' Business Continuity Plan (BCP) and Disaster Recovery Plan (DRP) through an end-to-end technology resiliency exercise, including physical failover testing of critical infrastructure and AI model provider dependencies.

1.2 Exercise Type

Physical end-to-end recovery test — combining tabletop scenario planning with executed failover procedures across production infrastructure.

1.3 Exercise Details

Field	Detail
Date	March 26, 2026
Time	10:00 AM WET

Duration	4 hours
Location	Virtual (Google Meet) and Lisbon Studio
Facilitator	Artur Ventura, CTO & CISO
Participants	Engineering Team, Product Team, Customer Support, Executive Leadership
Materials	BCP v1.0, DRP v1.1, Business Impact Analysis, Communication Plan, Contact Lists

2. Scenario Description

Scenario: Simultaneous Render Platform Outage and LLM Provider Degradation

Trigger event: Render.com experiences a major regional outage affecting all US-West services, including Synthetic Users' primary application hosting. Simultaneously, OpenAI's API experiences severe degradation with >50% request failures and 30-second+ latency on remaining requests.

Impact assessment:

- Complete loss of primary application hosting (Render)
 - Degraded AI/GenAI processing capability (OpenAI)
 - Client-facing SaaS platform unavailable
 - Active customer sessions interrupted
 - JPMC and other enterprise client SLAs at risk
-

3. Exercise Phases

Phase 1: Detection & Assessment (30 minutes)

Objective: Validate monitoring and alerting systems detect the outage within acceptable timeframes.

Step	Action	Owner	Target Time
1.1	Confirm Render outage detected via monitoring stack (Cloudflare, PaperTrail, Axiom)	Engineering Lead	< 5 min
1.2	Confirm OpenAI degradation detected via API error rate monitoring	Engineering Lead	< 5 min
1.3	Activate incident response team via communication plan	CTO	< 10 min
1.4	Assess outage scope and estimated duration (contact Render status page, OpenAI status)	Engineering Lead	< 15 min
1.5	Notify executive leadership and trigger BCP activation	CTO	< 20 min
1.6	Send initial client notification (enterprise clients including JPMC)	Customer Support	< 30 min

Result: Detection and assessment completed within 25 minutes. All monitoring alerts fired correctly. Incident response team assembled within 8 minutes via Google Meet.

Phase 2: Infrastructure Failover — Render to AWS (1.5 hours)

Objective: Execute physical failover of application hosting from Render to pre-configured AWS environment. Validate RTO of 4 hours.

Step	Action	Owner	Target Time
2.1	Activate pre-configured AWS ECS environment	Engineering Lead	< 15 min
2.2	Verify AWS RDS database accessibility and data integrity (RPO: 1 hour)	Engineering Lead	< 30 min
2.3	Deploy latest application container images to AWS ECS	Engineering	< 45 min
2.4	Update DNS records (Cloudflare) to point to AWS environment	Engineering Lead	< 50 min
2.5	Run end-to-end smoke tests on AWS-hosted application	Engineering	< 60 min
2.6	Verify client-facing functionality (login, data access, AI features)	Product Team	< 75 min
2.7	Confirm TLS certificates valid on AWS endpoint	Engineering	< 80 min
2.8	Monitor application stability on AWS for 10 minutes	Engineering	< 90 min

Result: Application successfully failovered to AWS. Total failover time: **3 hours 12 minutes** (within 4-hour RTO). Data integrity confirmed — last backup was 42 minutes old (within 1-hour RPO). All client-facing functionality operational on AWS.

Phase 3: AI Provider Failover — OpenAI to Anthropic (45 minutes)

Objective: Execute switch of AI/GenAI processing from OpenAI to Anthropic. Validate RTO of 2 hours.

Step	Action	Owner	Target Time
3.1	Activate Anthropic API configuration (LLM Shuffle feature flag)	Engineering Lead	< 5 min

3.2	Update model routing to Anthropic Claude models	Engineering	< 10 min
3.3	Run AI/GenAI functional tests (persona generation, synthetic user interviews)	Product Team	< 25 min
3.4	Validate output quality and response times against baseline	Product Team	< 35 min
3.5	Confirm no JPMC data leakage during provider switch	Engineering Lead	< 40 min
3.6	Monitor AI processing stability for 5 minutes	Engineering	< 45 min

Result: AI provider switch completed in **38 minutes** (within 2-hour RTO). All synthetic user generation and interview features operational on Anthropic Claude. Output quality validated against baseline — within acceptable parameters.

Phase 4: Communication & Client Management (concurrent)

Step	Action	Owner	Status
4.1	Post initial status page update (status.syntheticusers.com)	Customer Support	Completed
4.2	Send enterprise client notification (JPMC, others) per SLA requirements	Customer Support	Completed
4.3	Send 1-hour progress update to enterprise clients	Customer Support	Completed
4.4	Send resolution notification once failover complete	Customer Support	Completed
4.5	Confirm JPMC 72-hour incident notification obligation met	CTO	Confirmed

Phase 5: Debrief & Evaluation (45 minutes)

5.1 Results Summary

Metric	Target	Actual	Status
Detection time	< 15 min	5 min	PASS
Incident team assembly	< 20 min	8 min	PASS
Infrastructure failover (Render → AWS)	< 4 hours (RTO)	3h 12m	PASS
Data integrity (RPO)	< 1 hour	42 min	PASS
AI provider switch (OpenAI → Anthropic)	< 2 hours	38 min	PASS
Client notification	< 30 min	22 min	PASS
End-to-end service restoration	< 4 hours	3h 50m	PASS

5.2 Findings & Lessons Learned

- DNS propagation delay** — DNS update took longer than expected (~15 minutes for full propagation). **Action:** Pre-configure lower TTL values on critical DNS records.
- AWS environment drift** — Minor configuration differences between Render and AWS environments required manual fixes during failover. **Action:** Implement monthly sync checks between environments.
- LLM Shuffle worked as designed** — The model routing feature flag enabled seamless provider switching without code changes. No action needed.
- Client communication templates effective** — Pre-drafted templates reduced notification time significantly. **Action:** Add JPMC-specific template with SLA references.
- Data integrity validated** — Automated backup verification confirmed RPO compliance. No data loss detected.

5.3 Remediation Actions

Finding	Action	Owner	Due Date
---------	--------	-------	----------

DNS propagation delay	Reduce TTL to 60s on critical records	Engineering Lead	April 15, 2026
AWS environment drift	Implement monthly environment sync checks	Engineering	April 30, 2026
JPMC-specific notification	Create JPMC incident notification template	Customer Support	April 10, 2026

4. Conclusion

The March 2026 Business Continuity & Technology Resiliency Exercise successfully validated Synthetic Users' ability to recover from a simultaneous infrastructure and AI provider failure. All Recovery Time Objectives and Recovery Point Objectives were met. The exercise confirmed end-to-end recovery capability including physical failover, data integrity verification, AI provider switching, and client communication procedures.

5. Sign-Off

Role	Name	Date
CTO & CISO	Artur Ventura	March 26, 2026
Engineering Lead	[Name]	March 26, 2026

Next scheduled exercise: September 2026 (semi-annual)