Synthetic Users AI/GenAI Decommissioning Policy

Document ID: CRA-4.2.1-AIDC-001

Version: 1.0

Effective Date: March 25, 2026

Last Updated: March 25, 2026

Owner: CTO — Artur Ventura

Approved By: CEO — Kwame Ferreira

Classification: Internal – Confidential

CRA Control: CRA 4.2.1

---

1. Purpose

This policy defines the procedures for retiring, decommissioning, or replacing AI/GenAI systems, components, and third-party model integrations at Synthetic Users. It ensures that decommissioning activities are performed in a controlled, secure, and auditable manner — with particular attention to data handling, access revocation, and evidence of asset retirement.

This policy supplements the SDLC AI/GenAI Addendum and applies specifically to the end-of-life phase of the AI/GenAI model lifecycle.

---

2. Scope

This policy applies to:

- All third-party LLM provider integrations (e.g., OpenAI, Anthropic, Google Gemini, Mistral) accessed via the LLM Shuffle framework

- All prompt templates and agent workflow configurations managed in the platform
- All RAG pipeline components including vector indices, embedding models, and document ingestion workers
- All AI/GenAI features removed from production
- Internal AI-powered tools that process company or customer data

---

3. Decommissioning Triggers

An AI/GenAI decommissioning event is initiated when:

Trigger	Description
Provider deprecation	LLM provider announces end-of-life for a model version used in production
Security vulnerability	A material security or privacy issue is identified with a model or provider
Contract termination	Commercial relationship with an LLM provider is terminated
Strategic replacement	A model or provider is replaced by a superior alternative
Feature retirement	An AI/GenAI feature is removed from the platform
Regulatory requirement	A regulatory obligation requires retirement of a specific AI capability

---

4. Decommissioning Lifecycle

4.1 Planning

- Engineering Lead documents the decommissioning scope, timeline, and impact assessment

- Affected systems, data stores, API keys, and prompt templates are identified
- A migration path or replacement is confirmed before decommissioning begins (except in emergency security events)
- Timeline is set with a minimum 30-day notice period for planned decommissions

4.2 Migration & Transition

- Replacement model or feature is validated per the SDLC AI/GenAI Addendum model validation checklist before the deprecated system is retired
- Traffic is migrated progressively (canary → full cutover) to the replacement
- All prompt templates are updated and re-tested against the replacement model
- Customer-facing communications are issued if the change affects visible product behaviour

4.3 Access Revocation

- API keys and credentials for the decommissioned provider or model are rotated and retired immediately upon cutover
- Service accounts and OAuth tokens associated with the decommissioned system are revoked
- CI/CD pipeline references are removed to prevent accidental re-deployment
- Access revocation is confirmed and logged in the change management record

4.4 Data Removal

Synthetic Users does not retain customer data on behalf of LLM providers (providers operate under DPAs prohibiting data retention). Upon decommissioning a provider, the following data removal actions are performed:

Data Category	Removal Action	Evidence Required
API request/response logs at provider	Confirm provider's data retention policy; request deletion if retention occurred	Written confirmation from provider or DPA clause reference

RAG vector index chunks (tenant-specific)	Purge tenant vector index if the embedding model is also being decommissioned	Deletion log from vector store
Prompt template versions	Archive in version-controlled repository; mark as deprecated	Git tag + changelog entry
Cached inference results	Flush any application-level caches referencing the decommissioned model	Deployment log confirming cache clear
Internal test data	Delete test prompts and outputs generated during integration testing	Developer confirmation in change record

4.5 Verification

- Engineering Lead confirms all access credentials are revoked
- CTO reviews and signs off on data removal evidence
- Remaining system dependencies on the decommissioned model are verified as zero (automated dependency scan)
- A post-decommission test is run to confirm the decommissioned endpoint is no longer called in production

4.6 Documentation & Closure

- Decommissioning record is created including: scope, timeline, data removal evidence, access revocation confirmation, and sign-offs
- Data flow diagram is updated to remove the decommissioned model or provider
- Subprocessors list is updated if a provider is fully exited
- Change management record is closed and archived per the Information Governance & Records Management Standard

---

5. Separation of Duties

Role	Responsibility
CTO — Artur Ventura	Initiates and approves decommissioning plan; signs off on data removal evidence
Engineering Lead	Executes migration, access revocation, and data removal; maintains change record
Legal	Reviews provider DPA to confirm data deletion obligations and obtain written confirmation
CEO — Kwame Ferreira	Notified of any decommissioning that affects commercial relationships or contractual obligations

No single person may both execute and verify data removal steps — the CTO must independently verify evidence provided by the Engineering Lead.

6. Emergency Decommissioning

In the event of an active security incident requiring immediate decommissioning of a model or provider:

1. CTO declares emergency decommissioning — verbal or written authorisation is sufficient to initiate
2. Engineering Lead immediately revokes API keys and disables the affected integration
3. LLM Shuffle failover to an alternate provider is activated
4. Standard decommissioning documentation is completed within 48 hours of the emergency action
5. Incident response procedures are triggered in parallel per the Incident Response Plan

7. Periodic Asset Recertification

In addition to event-triggered decommissioning, all AI/GenAI assets are subject to annual recertification:

Asset Type	Recertification Activity	Frequency
Active LLM provider integrations	Confirm DPA is current; re-validate security posture	Annual
Prompt templates	Review for accuracy, safety, and alignment with current model capabilities	Annual
RAG index sources	Confirm source documents are current; remove stale or unauthorized content	Annual
LLM provider API keys	Rotate and re-issue	Annual (minimum)
AI/GenAI subprocessor list	Confirm all active providers are listed; remove decommissioned providers	Annual

Assets that fail recertification are treated as decommissioning events.

8. Data Standards by Asset Type

8.1 Cloud/AI Provider Data

- Customer data is protected by DPAs with all providers; providers are contractually prohibited from retaining API request data
- Upon provider exit, written confirmation of data deletion is obtained and retained for 5 years
- If written confirmation cannot be obtained, the risk is escalated to the CTO and documented in the decommissioning record

8.2 Internal Platform Data

- Vector index data (RAG) is deleted from the vector store upon decommissioning of the associated embedding model
- Application-level caches are flushed as part of the deployment process
- Database records referencing decommissioned models are updated to reflect retired status (not deleted, to preserve audit trail)

8.3 Exception Handling

If data removal cannot be completed due to legal hold, regulatory requirement, or technical constraint:

- The exception is documented in the decommissioning record with the specific reason
- A risk assessment is performed and approved by the CTO
- A remediation plan and target completion date are established
- The exception is reviewed at the next quarterly security review

---

9. Review

This policy is reviewed annually or following any major AI/GenAI decommissioning event. Updates are approved by the CTO.

---

10. Related Documents

- SDLC AI/GenAI Addendum
- AI/GenAI Algorithm Design Document
- Information Governance & Records Management Standard
- Third-Party Risk Management Policy
- Incident Response Plan

- Subprocessors & Data Flow

Synthetic Users, Inc. — 4223 Glencoe Ave, Suite C215-523, Marina del Rey CA 90292