Change Management Policy

Synthetic Users

Version: 1.2

Effective Date: November 2025

Last Updated: March 25, 2026

Document Owner: CTO — Artur Ventura

Review Frequency: Annually or upon significant system or regulatory change

Classification: Internal – Confidential

Change History

Version	Date	Author	Changes
1.1	November 2025	CTO / Security Lead	Prior release
1.2	March 25, 2026	Artur Ventura, CTO	Added explicit AI/GenAI change classification and governance section (Section 10); updated scope to explicitly reference AI/GenAI systems and model configuration; updated related documents. Per JPMC SCA CRA 26.1.8.

---

1. Purpose

This policy defines the controls and processes used by Synthetic Users to manage changes to production systems in a controlled, auditable, and secure manner, while

minimizing risk to availability, security, and customer data.

---

2. Scope

This policy applies to:

- Application code
- Infrastructure and cloud configuration (AWS, Render, managed services)
- Security configurations (TLS/SSL, IAM, network controls)
- **AI/GenAI systems** — including LLM orchestration logic, model provider configuration, system prompt changes, RAG pipeline updates, Persona Engine logic, and LLM Shuffle routing rules
- Third-party service integrations

The policy applies to all employees and authorized contractors with change privileges.

---

3. Change Classification

Changes are categorized as follows:

a. Standard Changes

Low-risk, repeatable changes with predefined procedures (e.g. dependency updates, configuration tuning).

- Pre-approved
- Logged automatically via version control and deployment tooling

b. Normal Changes

Planned changes that may impact production behavior.

- Require peer review and approval

- Tested prior to deployment
- Deployed via controlled CI/CD pipelines

c. Emergency Changes

Changes required to address active security issues, incidents, or service degradation.

- May bypass standard approval timelines
- Must be documented and reviewed retrospectively
- Common triggers: vulnerability remediation, security misconfiguration, platform incidents, LLM provider security events

---

4. Change Process

a. Request & Documentation

All changes are tracked through:

- Version control (Git)
- Pull requests with documented intent, scope, and risk
- Issue tracking where applicable

Each change includes:

- Description of the change
- Risk assessment (availability, security, data impact)
- Rollback considerations

b. Review & Approval

- All non-emergency changes require peer review prior to merge
- Security-impacting changes receive additional scrutiny
- Approval is enforced via repository protections

c. Testing

Changes are validated using:

- Automated tests where applicable
- Staging or controlled pre-production environments
- Targeted validation for security and configuration changes (e.g. TLS, access controls)

d. Deployment

- Deployments are executed via CI/CD pipelines
- Production access is restricted and role-based
- Infrastructure changes are managed via platform-native controls

e. Rollback

- Rollback procedures are defined per change
- Versioned deployments allow rapid reversion if required

---

5. Emergency Changes

Emergency changes:

- Are logged and traceable
- Are reviewed post-implementation
- Are limited to resolving the immediate issue

Post-incident reviews document:

- Root cause
- Actions taken
- Preventive controls

---

6. Security & Compliance Alignment

- Security-related changes follow the Vulnerability Management Plan
- Low, medium, and high-severity findings are remediated within defined SLAs
- TLS/SSL, IAM, and network changes are validated post-deployment
- Changes are auditable and retained for compliance purposes

7. Segregation of Duties

- No single individual can unilaterally approve and deploy high-risk changes
- Access to production systems is limited to authorized personnel
- Access is reviewed periodically

8. Monitoring & Logging

- System behavior and errors are continuously monitored
- Deployment activity is logged
- Security and operational alerts trigger investigation and remediation

9. Review & Maintenance

This policy is reviewed at least annually or upon significant changes to:

- Infrastructure
- Regulatory requirements
- Security posture
- AI/GenAI system architecture or model provider relationships

10. AI/GenAI Change Governance

AI/GenAI systems at Synthetic Users (LLM orchestration, Persona Engine, RAG pipeline, LLM Shuffle) introduce change categories that require additional governance controls beyond standard software change management. All AI/GenAI changes are in scope for this policy and subject to the same classification, review, testing, and documentation requirements.

10.1 AI/GenAI Change Categories

Change Category	Examples	Classification
Model Provider Switch	Adding, removing, or reprioritizing LLM providers in LLM Shuffle	Normal
System Prompt Change	Modifying Persona Engine system prompts or context injection templates	Normal — requires CTO sign-off
RAG Configuration Change	Modifying retrieval scope, embedding model, chunking strategy, or tenant isolation controls	Normal
Model Version Upgrade	Upgrading to a new model version from an existing provider	Normal — requires regression testing
LLM Provider Credential Rotation	Rotating API keys for OpenAI, Anthropic, Google, or other LLM providers	Standard
Content Filtering Adjustment	Modifying output filtering rules or safety thresholds	Normal — requires security review
Emergency Model Fallback	Switching to a fallback provider due to provider outage or security event	Emergency

10.2 AI/GenAI-Specific Review Requirements

Changes to AI/GenAI systems require the following review steps in addition to standard peer review:

- **Behavioral regression testing** — confirm that model output quality and accuracy meet baselines after the change
- **Tenant isolation verification** — confirm that RAG retrieval and context injection do not cross tenant boundaries following the change
- **Security impact assessment** — evaluate whether the change affects prompt injection risk, data leakage risk, or content filtering effectiveness
- **Provider DPA compliance check** — for changes that add or modify LLM provider relationships, confirm the provider's Data Processing Agreement is in place and current

10.3 Documentation and Auditability

All AI/GenAI changes must be documented with:

- The specific AI/GenAI component affected (orchestration layer, Persona Engine, RAG pipeline, LLM Shuffle, content filtering)
- The model provider(s) affected, if applicable
- Pre- and post-change behavioral testing results
- CTO approval record for system prompt and provider changes

Records are retained per the Information Governance & Records Management Standard.

---

11. Related Documents

- Secure SDLC Policy — Development lifecycle controls governing code changes
- SDLC AI/GenAI Addendum — AI/GenAI model lifecycle and adversarial testing requirements
- Incident Response Plan — Emergency change triggers and post-incident review

- Third-Party Risk Management Policy — LLM provider risk classification and vendor change controls
- Information Governance & Records Management Standard — Change record retention requirements