



Strategy + Creative + Tech

SoTech DAaaS — Fair Usage Policy (FUP)

Version 1.0

October 2025

1. Purpose

This Fair Usage Policy (“Policy”) ensures that all SoTech DAaaS users benefit from a stable, secure, and high-performance service.

By design, SoTech DAaaS is built to scale; however, to maintain performance and fairness across all customers, usage is continually monitored and assessed against typical patterns for each subscription tier.

2. Scope

This Policy applies to all usage of the SoTech Digital Agent as a Service platform, including:

- Chat interactions and message processing
- Website and document ingestion (scraping, uploads, API feeds)
- Vector generation, search, and embedding operations
- API requests and background processing tasks

It applies to all users, agents, and tenants within the SoTech multi-tenant environment.

3. Our Commitment

SoTech aims to provide unrestricted access to platform capabilities without imposing hard limits on:

- Number or size of uploaded documents
- Amount of content scraped or processed
- Volume of conversations or queries

We trust our users to operate responsibly. However, unrestricted access relies on shared responsibility — excessive, automated, or disproportionate use can impact service quality for others.

4. Monitoring & Performance Safeguards

To protect service availability and quality, SoTech continuously monitors platform activity using automated metrics and performance analytics.

This includes monitoring:

- Request frequency and API call volumes
- Scraping and embedding activity rates
- Vector storage growth and query patterns
- Unusual or automated behaviour inconsistent with typical user activity

Monitoring is performed only for performance and security purposes, and not for analysing customer data content.

5. Embedding and Content Ingestion

To maintain fair system performance and predictable API costs, SoTech monitors embedding activity across all users.

While there are no strict limits on the number of documents or URLs that can be uploaded, content must remain within reasonable and optimised boundaries.

5.1 Word and Token Volume Guidelines

- Each individual document, web page, or content item ingested should ideally not exceed 50,000 words (approx. 75,000 tokens).
- Uploads or ingestions above 100,000 words (≈150,000 tokens) may be split, rejected, or queued automatically to maintain performance.
- For websites, the automated scraper applies intelligent limits per page and may throttle ingestion if volume or page count becomes excessive.
- Large or repetitive content sets (e.g. log data, code dumps, or raw PDFs) are not suitable for embedding and may be flagged for review.

5.2 Content Optimisation

Users are encouraged to:

- Upload clean, relevant text (not images, scans, or redundant content).
- Avoid uploading entire websites in one batch — focus on key pages or knowledge sections.
- Use shorter, focused documents to improve embedding accuracy and query speed.

5.3 Oversized Embedding Management

If a user's ingestion volume or token count consistently exceeds the typical range for their subscription tier, SoTech may:

1. Notify the user with optimisation recommendations.
 2. Apply automated batching or rate limiting to embedding tasks.
 3. Offer a higher-capacity plan or dedicated vector storage option for sustained large-scale usage.
-

6. When Fair Usage Applies

If a user's activity:

- Significantly exceeds typical usage for their subscription tier, or
- Adversely affects platform stability, performance, or other users,

SoTech may take proportionate action, such as:

1. Notifying the user and providing a summary of the observed usage.
2. Recommending adjustments to ingestion or automation patterns.
3. Introducing temporary throttling, queue delays, or rate limits.
4. Discussing migration to a higher or dedicated plan if sustained heavy usage is expected.

These actions are always intended to maintain platform reliability for all users.

7. Prohibited Use

Users must not:

- Use the DAaaS platform for illegal, unethical, or abusive purposes.
- Circumvent system safeguards, or attempt to interfere with monitoring or quota controls.
- Employ automated or scripted ingestion that overloads scraping or embedding microservices.
- Use the platform to replicate, host, or resell AI outputs or training data outside their own organisation.

8. Transparency & Communication

SoTech will always:

- Communicate clearly before applying any fair usage restrictions.
- Provide data-driven explanations if intervention is required.
- Offer solutions, such as workload optimisation or upgraded capacity, before applying limits or suspensions.

9. Updates

This Policy may be updated as SoTech scales and new features are introduced. The latest version will always be available at: **<https://sotechnology.co.uk/fair-usage-policy>**

10. Contact

For questions about this Policy or to discuss usage concerns, please contact:
info@sotechnology.co.uk