



Quickstart

Generate AI video, images, and chat through an OpenAI-compatible API. If you already use the OpenAI SDK, change one line — the `base_url` .

↓ Download API reference (PDF)

1. Create an `sk-` key in the console [API Keys](#)
2. Point `base_url` at `https://api.perseworks.hk/v1`
3. Submit a video job, then poll for the result

[cURL](#)[Python](#)[JavaScript](#)

Generate a video

Video is async — submit returns a task id; poll it until it succeeds.

```
curl https://api.perseworks.hk/v1/video/generations \
-H "Authorization: Bearer sk-..." \
-H "Content-Type: application/json" \
-d '{
  "model": "seedance-2.0",
  "prompt": "a cat surfing a wave, cinematic",
  "resolution": "1080p",
  "durationSec": 5
}'
```

[Copy](#)

`resolution` : 720p · 1080p · 4K (1080p and 4K on Seedance 2.0 only; Fast/Mini are 720p). `aspectRatio` : 16:9 · 4:3 · 1:1 · 3:4 · 9:16 · 21:9 · adaptive. `durationSec` : up to 15. `image` : optional reference-image URL for image-to-video.

Poll for the result

```
curl https://api.perseworks.hk/v1/video/generations/vid_... \
-H "Authorization: Bearer sk-..."
```

[Copy](#)

```
# → { "status": "running" }           (keep polling)
# → { "status": "succeeded", "video_url": "https:// ...", "duration": 5 }
```

Generate an image

Images are synchronous — one call returns the result URL(s). Powered by Seedream 5.0 Pro.

```
curl https://api.perseworks.hk/v1/images/generations \
-H "Authorization: Bearer sk-..." \
-H "Content-Type: application/json" \
-d '{
  "model": "seedream-5.0",
  "prompt": "a red panda barista, cinematic lighting",
  "size": "2K"
}'
```

Copy

`size` : 1K · 2K · 3K · 4K, or an explicit `WxH` (e.g. 1024×1024). `image` : optional reference image(s) — URL or base64, up to 10 — for image-to-image / multi-reference. `n` : 1–4 images. `seed` : reproducible.

Response — the generated image URL(s):

```
{
  "model": "seedream-5.0",
  "created": 1785000000,
  "data": [ { "url": "https:// ... /out.png" } ],
  "usage": { "images": 1 }
}
```

Chat (streaming)

```
curl https://api.perseworks.hk/v1/chat/completions \
-H "Authorization: Bearer sk-..." \
-H "Content-Type: application/json" \
-d '{
  "model": "seed-2.0-mini",
  "messages": [{"role": "user", "content": "Hello!"}],
  "stream": true
}'
```

Copy

Models

Pass the model id in your request. The live list is always `GET /v1/models`.

seedance-2.0

video

720p / 1080p / 4K · text + image to video · audio

<code>seedance-2.0-fast</code>	video	720p · faster generation
<code>seedance-2.0-mini</code>	video	720p · lowest cost
<code>seedream-5.0</code>	image	Seedream 5.0 Pro · text-to-image + editing · 1K or 2K · flat rate
<code>seed-2.1-pro</code>	chat	strongest chat model · 262K context
<code>seed-2.0-mini</code>	chat	fast + low cost · 262K context

Embeddings and audio are on the roadmap.

Endpoints

POST	<code>/v1/video/generations</code>	Submit a Seedance video job
GET	<code>/v1/video/generations/:id</code>	Poll a video job for its result
GET	<code>/v1/models</code>	List available models
POST	<code>/v1/chat/completions</code>	Chat completion (supports stream: true)
POST	<code>/v1/images/generations</code>	Generate Seedream images

Embeddings and audio endpoints are on the roadmap.

Errors & limits

Every error returns the same JSON shape alongside an HTTP status code:

```
{ "error": { "code": "invalid_request", "message": "a human-readable reason" } }
```

400 · <code>invalid_request</code>	Malformed body or a missing required field
400 · <code>content_blocked</code>	Prompt failed content moderation (never charged)
401 · <code>invalid_api_key</code>	Missing, unknown, or revoked API key
402 · <code>insufficient_points</code>	Not enough points to reserve the request
403 · <code>model_not_allowed</code>	This key is not scoped to that model
404 · <code>not_found</code>	Task not found, or not owned by this key
413 · <code>payload_too_large</code>	Request body exceeds the size cap

429 · rate_limited

Per-key rate limit exceeded — back off and retry

5xx · upstream_error / internal_error

Server or vendor error — the charge is refunded

Billing — requests are prepaid in points and pre-charged, then settled by actual usage (a failed request is refunded). Live prices: `GET /v1/models`.

Rate limits — applied per API key; exceeding them returns `429 rate_limited`. Back off and retry.

Idempotency — send an `Idempotency-Key` header on a video submit so a retry returns the same task instead of starting a new job.

Perseworks provides access to Seedance 2.0 and Doubao models as an authorized reseller, under licence, for international developers. "Seedance", "Doubao" and "火山方舟 (Volcano Ark)" are trademarks of Beijing Volcano Engine Technology Co., Ltd. / ByteDance. Perseworks is operated by Verve Group Inc., a company incorporated in Ontario, Canada.