

Contents

FAQ Kit AI Generator	1
1. Overview	1
Key features	1
Compatibility	2
2. Installation	2
Requirements	2
Via Composer	2
Verify	2
3. Configuration	3
3.1 Model and spend cap – LLM Provider	3
3.2 Generation settings – FAQ Kit	4
4. Generating FAQs	4
4.1 Where you start from	4
4.2 What the candidates are based on	5
4.3 Review and apply	5
5. Troubleshooting	6
6. Uninstallation	7
7. Support & Changelog	7

FAQ Kit AI Generator

User Guide

Version: 1.0.0

Compatibility: Adobe Commerce and Magento Open Source 2.4.7 – 2.4.9, PHP 8.1 – 8.5

Vendor: Webmaster Ramos

1. Overview

FAQ Kit AI Generator adds AI generation to FAQ Kit. From a product, a category, a FAQ widget or the FAQ Entries grid, it drafts candidate questions and answers from the item's catalog content or from text you supply, lists them on a review page, and saves only the ones you tick as FAQ Kit entries.

It is a companion package. FAQ Kit curates, renders and serves FAQs on its own; this package supplies the model call, through LLM Provider – the free, on-premise gateway that holds your provider key, the model choice and the spend caps. The package adds no screen of its own, no database table and no storefront output: it switches on the AI controls FAQ Kit already renders.

Key features

- Candidates grounded in a product's or category's catalog content
- Pasted text or an uploaded `.txt` / `.md` file as an extra source, or as the only source for store-wide and widget FAQs
- Four starting points: the FAQ Entries grid and the product, category and FAQ widget edit forms

- A review page where every candidate is a ticked row, with **Generate more**
- Only ticked candidates are saved; existing questions for the same target are left out
- Model and daily spend cap for FAQ generation chosen in LLM Provider
- Strict JSON output, one retry on a retryable provider error, a clear message when generation cannot run

Compatibility

Platform	Versions
Adobe Commerce	2.4.7 – 2.4.9
Magento Open Source	2.4.7 – 2.4.9
PHP	8.1, 8.2, 8.3, 8.4, 8.5
Requires Themes	FAQ Kit 1.0 or later, LLM Provider 1.0 or later Admin only – FAQs are shown by FAQ Kit on Luma and Hyvä Theme
Database	No tables added
External services	The model provider you configure in LLM Provider

2. Installation

Requirements

- FAQ Kit and LLM Provider. Composer installs both as requirements of this package if they are not installed yet.
- An API key for a model provider, entered in LLM Provider (section 3.1).

Via Composer

```
composer require webmaster-ramos/module-faq-kit-llm
bin/magento module:enable WebRamos_FaqKitLlm
bin/magento setup:upgrade
bin/magento setup:di:compile
bin/magento cache:flush
```

setup:upgrade also enables FAQ Kit and LLM Provider when they are new to the installation. setup:di:compile is required in production mode and harmless elsewhere.

Verify

```
bin/magento module:status WebRamos_FaqKitLlm WebRamos_FaqKit WebRamos_LlmProvider
```

All three should be enabled. In the admin, **Stores → Configuration → Webmaster Ramos → LLM Provider** now carries a **FAQ Kit** group under **Consumer Modules**, and the **Generate with AI** button on **Content → FAQ Kit → FAQ Entries** is no longer greyed out.

3. Configuration

3.1 Model and spend cap – LLM Provider

Stores → Configuration → Webmaster Ramos → LLM Provider

1. In the provider group you want to use, enter the **API Key** and save. The provider groups, **Test Connection** and the gateway's own settings are described in the LLM Provider user guide.
2. Open **Consumer Modules** → **FAQ Kit**.

The screenshot shows the 'Consumer Modules' configuration page for the 'FAQ Kit' module. On the left is a sidebar with icons for SALES, CATALOG, CUSTOMERS, MARKETING, LLM PROVIDER, CONTENT, REPORTS, STORES, and SYSTEM. The main content area is titled 'Consumer Modules' and contains a section for the 'FAQ Kit' module. Below the module name, it states: 'The model FAQ Kit may use, declared in its etc/llm_module.xml whitelist.' There are three configuration fields: 1. 'Model' (global) with a dropdown menu showing '-- Please Select --' and a note: 'Used for FAQ generation. Only whitelisted chat models are listed.' 2. 'Daily Spend Cap (USD)' (global) with a text input field and a note: 'Overrides the gateway-wide default cap for this module. Empty or 0 uses the gateway-wide Daily Spend Cap under General; there is no cap only when that is empty or 0 too.' 3. 'Allow Untested Models' (global) with a dropdown menu showing 'No' and a note: 'The module's declaration lists the models it was tested with. Set to Yes to also offer any other model in the registry that meets the module's stated requirements – typically one added to the Model Catalog after the module was released. Requirements such as structured output or a minimum context window are always enforced; everything else is on you to verify.'

Figure 1: Consumer Modules → FAQ Kit on a fresh install

Field	Meaning
Model	The chat model FAQ generation uses. The list offers the models this package was tested with – Claude Haiku 4.5 (recommended, listed first), Claude Opus 4.6, GPT-5 and GPT-5 mini – all of which support structured JSON output. Until a model is chosen, generation reports that the AI service is unavailable.
Daily Spend Cap (USD)	A daily budget for FAQ generation alone. When it is empty or 0, the gateway-wide cap under General applies, and when that is empty too there is no cap. The cap is a soft ceiling: the call that crosses it completes, the next one is refused until the next day.
Allow Untested Models	Added to every consumer group by LLM Provider. Yes also offers other models from the Model Catalog that meet this package's requirements – structured JSON output and a context of at least 16,000 tokens.

The cost of each call is estimated by LLM Provider from its price table and shown in its Usage Dashboard, under the module name `WebRamos_FaqKit`.

3.2 Generation settings – FAQ Kit

Stores → Configuration → Webmaster Ramos → FAQ Kit → General

These settings belong to FAQ Kit and are used by every generation run:

Field	Meaning
Enabled	FAQ Kit's master switch. When it is off for a store view, generation is off there too.
Generation Prompt	The base instruction sent before the source material. The shipped text asks for concise, useful FAQs grounded only in the source, with no invented facts, prices or policies; an empty field falls back to the same instruction.
Pairs per Generation	How many question and answer pairs a run asks for (5 by default), unless the review page sets another number.
Generation Language	<code>auto</code> follows the store view's locale; or a locale such as <code>en_US</code> .
Max Question Length / Max Answer Length	Characters per generated question and answer; 0 means no cap. Longer text is cut after generation.
Answer style	Neutral, Professional, Friendly, Concise, or a custom phrase of your own.

4. Generating FAQs

4.1 Where you start from

- **Content → FAQ Kit → FAQ Entries → Generate with AI** opens the review page unscoped: choose the kind – Global, Product, Category or Widget – the ID for the last three, and the store view.
- **Generate FAQs with AI** in the **FAQ** section of a product, category or FAQ widget edit form opens the same page already scoped to that item, with a heading naming it. On an item that has not been saved yet the button is greyed out – save the item first.




Figure 2: The FAQ section of a product form, with Generate FAQs with AI enabled

4.2 What the candidates are based on

- **Product:** its name, SKU, short description and description.
- **Category:** its name and description.
- **Source Text:** anything pasted into the review page.
- **Source File:** a .txt or .md file, up to 2 MB; text beyond 40,000 characters is not used.

For a product or a category, Source Text and Source File are optional and are added to the catalog content. The Global and Widget kinds have no catalog content, so pasted text or a file is required. Pasted text and uploaded files are sent with the request and never stored.

The source of each run is sent to the model provider chosen in section 3.1, and to no other destination.

4.3 Review and apply

1. Set **Store View**, **Pairs to Generate** (0 uses **Pairs per Generation**) and **Generation Language** if needed, then press **Generate**. A run can take up to a minute.
2. Every candidate is listed as a ticked row. Questions that already exist for the same target are left out.
3. **Generate more** adds further candidates below the listed ones, using the same source, and never repeats a question already on screen.
4. Untick what you do not want. **Enable saved entries** decides whether the saved FAQs are active straight away.
5. Press **Apply Selected**. Only the ticked rows are saved, through the same answer sanitising as the FAQ entry form, and assigned to the product, category or widget the page is scoped to. You return to the form you came from, or to the FAQ Entries grid for global FAQs.

- DASHBOARD
- SALES
- CATALOG
- CUSTOMERS
- MARKETING
- LLM PROVIDER
- CONTENT
- REPORTS
- STORES
- SYSTEM
- FIND PARTNERS & EXTENSIONS

Generate FAQs with AI

Generate FAQs for: Joust Duffle Bag

Set the parameters, generate candidates, then keep the ones you want. Nothing is written to your FAQs until you press Apply.

Source Text

Source File

[Choose File](#) No file chosen

Optional – added to the product or category content the page already uses. Text files only (.txt, .md).

Store View

All Store Views

Pairs to Generate

0

0 uses the configured default.

Generation Language

auto

auto follows the store view locale, or set a locale like en_US.

[Generate](#) [Back](#)

	Question	Answer
☒	What is the product name and SKU?	Joust Duffle Bag, SKU 24-MB01.
☒	What are the dimensions of the Joust Duffle Bag?	The bag measures 29 inches in length, 13 inches in width, and 11 inches in height.
☒	How can the Joust Duffle Bag be carried?	It features dual top handles and an adjustable shoulder strap for carrying.
☒	What type of closure does the Joust Duffle Bag have?	It has a full-length zipper closure.
☒	What items can the Joust Duffle Bag hold?	It can hold a basketball or soccer ball along with sneakers, with extra room to spare.
☒	Who is the Joust Duffle Bag intended for?	It is designed for athletes who need a bag for the gym, travel, or sports activities.
☒	Does the source specify the material of the Joust Duffle Bag?	The provided source does not include information about the bag's material or fabric.

☒ Enable saved entries

Adds new candidates below the ones listed, using "Pairs to Generate" above. The questions already listed are excluded.

[Generate more](#) [Apply Selected](#)

Copyright © 2026 Magento Commerce Inc. All rights reserved.

Magento ver. 2.4.8-p5
 Hyvä Theme Module ver. 1.4.5
[Privacy Policy](#) | [Account Activity](#) | [Report an Issue](#)

Figure 3: The review page for a product, with candidates generated from its catalog content

Nothing is written to your FAQs before **Apply Selected**. Saved entries are ordinary FAQ Kit entries: edit, disable or delete them like any other.

5. Troubleshooting

Symptom	Cause and resolution
The AI controls are greyed out with “AI generation is not installed”	This package is not installed or not enabled. Check <code>module:status</code> (section 2), then run <code>set up:di:compile</code> and <code>cache:flush</code> .

Symptom	Cause and resolution
The AI controls are greyed out with “FAQ Kit is turned off” “The AI service is unavailable right now”	FAQ Kit is disabled for that store view. Turn on Enabled under FAQ Kit → General. No model is chosen, the provider key is missing or wrong, the provider is unreachable, the daily spend cap is reached, or the gateway is switched off. The message links to the LLM Provider configuration; <code>var/log/system.log</code> carries the gateway’s error code.
“No source material”	A Global or Widget run with no pasted text or file, or a product or category with no content. Add text or upload a file.
“No suggestions were produced”	The model returned nothing usable. Run again, add source text, or change the pair count.
“All suggestions duplicated existing FAQs”	Every candidate matched a question the target already has. Add new source text.
“Unsupported file type”	Only <code>.txt</code> and <code>.md</code> files are read. Convert the document to plain text first.
The model you want is not in the list	It is not on this package’s tested list. Set Allow Untested Models to Yes ; the model must still support structured JSON output.

Errors are written to `var/log/system.log` and `var/log/exception.log`.

6. Uninstallation

Removing this package leaves FAQ Kit fully working: every FAQ stays, and the AI controls are shown greyed out with a hint.

Composer removes packages that nothing else requires. If FAQ Kit or LLM Provider were installed only as requirements of this package, first add `webmaster-ramos/module-faq-kit` (and, if other extensions use it, `webmaster-ramos/module-llm-provider`) to the `require` section of your project’s `composer.json`, so they stay installed.

```
bin/magento module:disable WebRamos_FaqKitLlm
composer remove webmaster-ramos/module-faq-kit-llm
bin/magento setup:upgrade
bin/magento setup:di:compile
bin/magento cache:flush
```

The model and spend cap saved under **Consumer Modules** → **FAQ Kit** stay in the configuration table and are used again if the package is reinstalled.

7. Support & Changelog

- **Author:** Webmaster Ramos

- **Email:** contact@webmaster-ramos.com
- **Website:** <https://webmaster-ramos.com>
- **Support:** <https://webmaster-ramos.com/support>
- **Changelog:** see `CHANGELOG.md` inside the package