

Verso

Product Description and Functional Specification

Document ID	IM.VERSO.INT.01
Document version	1.0
Date of issue	2026-10-06
Applies to	Verso 2026-09-25 (4bac47a)
Language	English
Publisher	iMobile

Contents

- 1. General information 3
 - 1.1. Name and release identifier 3
 - 1.2. Purpose 3
 - 1.3. Scope of use 3
 - 1.4. Product class 3
 - 1.5. Origin of the code base 3
- 2. Functions 3
 - 2.1. Creating a presentation 3
 - 2.2. Outline 4
 - 2.3. Templates and design 4
 - 2.4. Slide editor 4
 - 2.5. AI assistant 4
 - 2.6. Export and presenting 4
 - 2.7. Accounts and workspaces 4
 - 2.8. Admin console 4
 - 2.9. Credit accounting 5
 - 2.10. On-premises licensing 5
- 3. Architecture and components 5
- 4. Interfaces and data exchange 5
 - 4.1. User interface 5
 - 4.2. Programming interface 5
 - 4.3. Model providers and external services 6
 - 4.4. Data flows by deployment option 6
 - 4.5. Data formats 6
- 5. Editions and limitations 7
 - 5.1. Delivery options and editions 7
 - 5.2. Known limitations 7
- 6. Environment requirements 7

About this document

This document describes the purpose, functions, components and interfaces of Verso, a server-side web application that prepares editable presentations with generative AI models. It is written for the customer's specialists who assess the product before purchase and deployment.

It covers the release of 2026-09-25 in both delivery options: the vendor cloud service and the on-premises installation. Where the two differ, the text says so.

1. General information

1.1. Name and release identifier

The product name is Verso. Verso has no release numbering of its own: a release is identified by its date and the short identifier of the source revision. This document applies to release 2026-09-25 (revision 4bac47a). The web interface does not display a version number.

1.2. Purpose

Verso builds a presentation from a text prompt or uploaded documents: it drafts an outline, picks slide layouts from the selected template, and writes the text, charts, tables, images and icons. The result is edited in the browser and exported to a PowerPoint (PPTX) or PDF file.

1.3. Scope of use

Preparing business, reporting and training presentations in organisations. The on-premises installation is intended for cases where source material must not leave the customer's network.

1.4. Product class

Application software: a server-side web application that uses external or local generative AI models. No language model is part of the product.

1.5. Origin of the code base

Verso is built on the open-source Presenton project (Apache License 2.0), from a source snapshot of 2026-08-03 that the source tree labels 0.9.3-beta. iMobile's additions are credit accounting and payments, token usage accounting and the admin console overview, on-premises licensing, sign-in through the vendor site account, interface translation, extra templates, server-side generation jobs, and removal of usage analytics reporting. The licence verification module is iMobile's own closed-source component.

2. Functions

2.1. Creating a presentation

- Input: a text prompt and/or up to 8 attached files of up to 100 MB each.
- Attachment formats: PDF, TXT, Word documents (DOC, DOCX, DOCM, ODT, RTF), presentations (PPT, PPTX, PPTM, ODP), spreadsheets (XLS, XLSX, XLSM, ODS, CSV, TSV), images (JPG, PNG, GIF, BMP, TIFF, WEBP). Text is extracted on the server; scans and images go through text recognition.
- Parameters: number of slides (automatic or set, up to 50), presentation language (auto-detect or a choice from a list), tone, verbosity, extra instructions, table of contents, title slide.
- The "Use my text as-is" mode: the model only chooses the design and does not rewrite the user's text.

- Generation runs on the server as a background job. Closing the browser tab does not stop it, and progress can be watched by reopening the presentation.

2.2. Outline

Before the slides are created the program drafts an outline. The user edits the items (Markdown), reorders them, adds items, regenerates the outline from a changed prompt, or asks the AI assistant to revise it.

2.3. Templates and design

- Built-in templates (12): Momentum, Dynamic, Executive, General, Modern, Standard, Swift, Biology Mint, School Bright, School Classic, Secure Dark, Software Showcase. Their fonts are included in the image.
- Custom templates are created from the user's PPTX file; during processing a model analyses images of the slides.
- Themes: the colours and fonts of a presentation can be changed in the editor.

2.4. Slide editor

- Text editing; moving, resizing, grouping and layering objects on the canvas; undo and redo; keyboard shortcuts.
- Inserting text blocks, charts (bar, horizontal bar, line, pie, donut, area, scatter, radar and others), tables, images, shapes and ready-made template blocks.
- Adding, duplicating, reordering and deleting slides; speaker notes.
- Images: stock search, AI generation, upload of the user's own file (up to 5 MB), fit and focus point settings.

2.5. AI assistant

The assistant takes plain-language requests about the outline and the finished presentation: rewrite or shorten text, add a slide, chart or table, replace an image, write speaker notes. A file or selected slide objects can be attached to a request. Changes are shown for comparison and the original can be restored.

2.6. Export and presenting

- Export to PPTX (editable objects) and to PDF.
- Presentation mode in the browser, with speaker notes.

2.7. Accounts and workspaces

Sign-in with a user name and password. Roles: administrator and user. Each user has a separate workspace; presentations, templates and images of other users are not accessible.

2.8. Admin console

- Overview: number of users and presentations, token usage by model, running jobs, credit and licence status.
- Per-user table: credit balance and spend, number of presentations and slides, token usage, last activity.
- Account management: create, reset password, change role, delete.
- Configuration of the text model, image and web search providers.
- API and MCP access keys; promo codes for credits.

2.9. Credit accounting

In the cloud service generation is paid for with credits. Credits are reserved before a run starts and charged for the slides actually produced; the reservation is returned on failure. In this release a slide costs 1, 3 or 5 credits depending on the model class; a message to the assistant costs 1 credit plus the slide cost for each slide the assistant writes in full. The balance is topped up by buying a credit package or with a promo code; the ledger is visible to the user. On an on-premises installation credit charging is off by default.

2.10. On-premises licensing

An on-premises installation runs under a licence file signed by the vendor. The verification module checks the Ed25519 signature against a built-in public key, the product identifier, the expiry date and, if set, the installation binding. The check is local and does not contact the vendor. If the licence is missing, damaged or expired, starting presentation generation is blocked; presentations created earlier remain available for viewing, editing and export. The licence status is shown in the admin console.

3. Architecture and components

The product is shipped as a single Docker image. The container runs:

Component	Role
nginx	Front proxy, container port 80 (HTTP); serves files after a session check
Web interface (Next.js)	Application pages, slide editor
API server (FastAPI)	Generation, storage, accounts, credit and token accounting
MCP server	Access to generation over the Model Context Protocol
Document parsing	Text extraction from files, text recognition
Export	Renders slides in Chromium and produces PPTX and PDF
Licence verification module	On-premises only; shipped in compiled form

Data is kept in `/app_data`: the database (SQLite by default; external PostgreSQL or MySQL by configuration), the settings file, images, exported files, fonts, templates and the presentation memory store. The application server runs as a single worker process.

4. Interfaces and data exchange

4.1. User interface

A web interface over HTTP; channel protection (HTTPS) is provided by a reverse proxy in front of the container.

4.2. Programming interface

The product provides a REST API that exchanges JSON. Main operations:

Operation	Method and path
File upload	POST <code>/api/v1/ppt/files/upload</code>
Presentation generation	POST <code>/api/v1/ppt/presentation/generate</code>
Background generation	POST <code>/api/v1/ppt/presentation/generate/async</code>
Background generation status	GET <code>/api/v1/ppt/presentation/status/{id}</code>

Operation	Method and path
Webhook subscription	POST /api/v1/webhook/subscribe

- The result of generation is a PPTX or PDF file and a link for editing the presentation.
- Authentication uses an access key in the `Authorization: Bearer` header; keys are issued by an administrator. The OpenAPI description is available at `/docs` after sign-in.
- The MCP server at `/mcp` accepts the same access key; tools: generate a presentation, generate in the background, query status.
- Webhook notifications are sent to an external address when generation completes or fails.

4.3. Model providers and external services

Purpose	Supported options
Text model	OpenAI, Anthropic, Google (Gemini, Vertex AI), Azure OpenAI, Amazon Bedrock, DeepSeek, OpenRouter, Fireworks, Together, Cerebras, a LiteLLM gateway, any server with an OpenAI-compatible API; locally Ollama and LM Studio
Images	Pexels and Pixabay stock libraries; OpenAI and Google models; local ComfyUI and Open WebUI servers; a server with an OpenAI-compatible API; image generation can be switched off
Web search	The model provider's built-in search (OpenAI, Google, Anthropic), SearXNG, Tavily, Exa; off by default

4.4. Data flows by deployment option

Data	On-premises, local model	On-premises, external provider	Cloud service
Prompt, instructions, text of attached documents, outline, slide content, assistant requests	Stay on the customer's server	Sent to the text model provider chosen by the customer's administrator	Processed on the vendor's servers and sent to the model provider chosen by the vendor
Image descriptions and stock search queries	Not sent if images are off or a local image server is connected	Sent to the chosen image provider	Sent to the image provider chosen by the vendor
Slide images of a PPTX file when a custom template is created	Processed by the local model	Sent to the text model provider	Sent to the model provider chosen by the vendor
Accounts, presentations, files, token accounting	On the customer's server	On the customer's server	On the vendor's servers
Data passed to the vendor	None	None	All of the above

Usage analytics is not sent in this release: the analytics client is inactive without a project token, and the release sets none. An on-premises installation does not check for updates and does not contact the vendor's servers.

4.5. Data formats

Input: the attachment formats listed above; PPTX for custom templates; TTF and OTF font files. Output: PPTX, PDF.

5. Editions and limitations

5.1. Delivery options and editions

Option	Terms
Cloud service	Pay-as-you-go credits; model providers are chosen by the vendor
On-premises, Team	Annual licence, 25 users
On-premises, Business	Annual licence, 100 users
On-premises, Enterprise	Annual licence, no user limit

The functions of the program are the same in all on-premises editions. The edition, the number of users and the list of edition features are written into the licence file and shown in the admin console. The program checks that the licence is genuine and has not expired; it does not enforce the number of users or the edition features.

5.2. Known limitations

- Content is written by a model and may contain inaccuracies; it must be reviewed by a person.
- No multi-factor authentication, no SAML or OIDC sign-in, no LDAP integration; user passwords are reset by an administrator.
- Open registration should not be enabled on an on-premises installation; accounts are created by the administrator.
- No audit log of user or administrator actions.
- No encryption at rest; provider API keys are stored in clear text.
- One container and one worker process; a restart interrupts running generations, finished slides are kept.
- Fonts are not embedded in an exported PPTX.
- Deleting a presentation does not delete its exported files, images or memory entries.
- For an isolated network: building the image needs internet access; fonts from the theme catalogue are loaded by the browser from Google Fonts; the vendor has not audited the own network activity of third-party libraries in the image.
- The admin console and the custom template wizard are in English only.
- No accessibility assessment has been carried out.

6. Environment requirements

- On-premises: a host with Docker Engine and Docker Compose v2 that runs linux/amd64 containers; a port for the web interface (5000 in the package); a TLS reverse proxy for HTTPS access.
- A text model source: an API key of an external provider or a local model server reachable from the container. Hardware requirements for local models are determined by the chosen model.
- Workstation: a web browser.
- The vendor does not state minimum CPU, memory or disk figures for the application itself.

Document history

Document version	Date	Product version	Change
1.0	2026-10-06	2026-09-25 (4bac47a)	First issue