

Changelog

Contents

- [2026-07-19 - Resilient bootstrap downloads](#)
 - [Changed](#)
- [2026-07-17 - Canonical Workbench and shared CLI/GUI services](#)
 - [Changed](#)
 - [Verified](#)
- [2026-07-15 - Provider-neutral DocumentModel and office exports](#)
 - [Added](#)
 - [Changed](#)
 - [Clarified](#)
- [2026-07-06 - DEV Markdown PDF and GUI polish](#)
 - [Changed](#)
- [2026-07-01 - OCR Brick integration and UI cleanup](#)
 - [Added](#)
 - [Changed](#)
 - [Removed](#)
 - [Verified](#)
 - [Cleanup](#)

2026-07-19 - Resilient bootstrap downloads

Changed

- Restored the Workbench path/deletion helper chain so an unpinned external Source selected from history reaches the operation snapshot exactly like the native picker route.

- FZF and portable PowerShell installers now resolve and download deterministic release assets without consuming GitHub REST quota in the normal path, while preserving API metadata as a fallback when direct asset conventions change.
- Real-ESRGAN remains on its metadata-driven release resolver because its Windows portable asset does not follow the same simple latest-version contract as the bootstrap utilities.

2026-07-17 - Canonical Workbench and shared CLI/GUI services

Changed

- Transplanted the byte-identical canonical Workbench with the shared `Source`, `Add file...`, `Target`, `Reset`, `Delete`, and `List` vocabulary.
- Added protected single-file/folder pickers, process-tree cleanup, local-only GUI host enforcement and the sanitized native-window shutdown lifecycle.
- Routed GUI Workbench paths as per-operation Source/Target snapshots while preserving managed `input`/`output` roots for the backend mirror contract.
- Added `system_core/cli_operation.py`; both CLI/FZF launchers now execute the same manifest/service operations as NiceGUI for OCR Brick, Workbench resolvers and office builders.
- Replaced obsolete CLI-only OpenAI/Gemini OCR copies with the current Tesseract, Surya, Yandex, Mistral, xAI, Gemini and OpenAI OCR Brick routes.
- Hardened cleanup and release staging: `.git` is preserved, while user documents, reports, caches, local route history and Yandex credential metadata are excluded from portable archives.

Verified

- Added single-file mirror/Target synchronization tests, CLI/GUI dispatch parity checks and an offline Tesseract-to-DocumentModel/DOCX/JSON/Markdown integration smoke.
- Removed the unused legacy `extract_ai` service and generic model selector after confirming that no manifest, GUI, CLI or test referenced them.

2026-07-15 - Provider-neutral DocumentModel and office exports

Added

- Added the lossless `<name>.document` package as the canonical OCR layer, retaining exact sources/page rasters, hashes, all OCR candidates, coordinates, structured regions, table spans, confidence, verification and manual-correction slots.
- Added local DOCX, Searchable PDF, XLSX, ODT, Markdown, HTML, OCR JSON and full ZIP exporters. They do not rerun paid OCR and do not require LibreOffice.
- Added independent checkbox-chip output selection backed by the canonical DocumentModel.
- Added automatic local Tesseract coordinate generation when a cloud OCR result is exported as Searchable PDF.
- Capped adaptive DOCX page size at A3.
- Added content-aware column balancing for DOCX and XLSX: compact numeric/year columns yield space to names, addresses and other long text.
- Added Mistral second-pass GUI controls: none, Tesseract, Yandex, or Yandex plus Tesseract, with suspicious/all-page scope and optional text fusion.

Changed

- Output formats now use a balanced 4-by-2 checkbox-chip grid: four office formats and four developer formats. Full archive ZIP remains supported by the backend but is no longer exposed in the primary OCR GUI.
- Removed the duplicated `Output formats` and `Scan cleanup` field labels while retaining their shaded explanatory hint blocks; the section headings `Deliverables` and `Preprocessing` remain the only upper headings.
- Restored visual separation for deliverable chips by placing the complete 4-by-2 checkbox grid on the same subdued dark container treatment used by segmented button groups.
- Expanded the two OCR output-contract buttons into equal full-row columns for better visual balance.
- Added conservative Mistral/Yandex fusion. Table cells use verified physical geometry, provider-native coordinates or unambiguous native Tesseract word anchors; unverified grids never authorize replacement. Numeric disagreements and ambiguous coordinates remain review items.
- Added a physical ruled-table model: long raster lines define rows/columns, the OCR numbering row verifies the column count, header merges are recovered from missing separators, and all

provider text is assigned inside exact cells.

- Added Yandex candidates, provider verification and per-item fusion decisions to DocumentModel 1.1 and the XLSX audit sheet.
- Fixed the paid OCR root filter so Mistral OCR 4 is actually visible in the GUI.
- Kept the OCR engine selector in one horizontal row with a single **Engine** / **Движок** heading. Engine buttons now divide the full section width evenly: two local engines use two equal columns, and five API engines use five equal columns.
- Promoted the shared **Deliverables** / **Готовые файлы** block directly below the engine selector for both local and paid OCR. Preset buttons were removed; formats are independent checkboxes and DOCX is the working default. The internal DocumentModel is always preserved.
- Added a confirmed maintenance action **Clear DocumentModel artifacts** / **Очистить DocumentModel**. It removes regenerable ***.document**, ***.document.json** and ***.verification.json** artifacts from managed output/report/workspace while preserving DOCX, XLSX, Markdown, PDF, images and sources.
- Removed the obsolete **Build office output** card from the main GUI. The old Markdown compilers remain available under **Project tools -> Rebuild legacy Markdown** for compatibility with archived Markdown workflows.
- Reworked OCR deliverables into a balanced 4-by-2 grid of rounded rectangular checkbox chips: four equal amber office columns and four equal sea-green dev/machine columns. The chip outline is intentionally very subtle, while the square checkbox keeps a clearer but slightly muted outline; narrow layouts fall back to two columns.
- Flattened the **Tables** window into three expanded operation blocks with their own parameters and Run buttons; coordinate extraction, Markdown table inspection and XLSX build no longer require child navigation.
- Refreshed the Russian and English documentation stack around the current DocumentModel-first workflow, direct 4-by-2 deliverables grid, conservative table geometry, maintenance semantics and legacy-only Markdown rebuild path.

Clarified

- "100% completeness" means no loss of source data; it is not a claim of 100% OCR accuracy.

2026-07-06 - DEV Markdown PDF and GUI polish

Changed

- Changed the DEV Markdown PDF default output to `docs/PDF` in the GUI, service fallback and standalone renderer.
- Reworked DEV Markdown PDF source mode into compact top-level buttons: `Источник`, `input/output`, `Документация`, `PDF-пары`, `Устаревшие PDF`.
- Kept `PDF рядом с MD` as an explicit alternative: every folder with Markdown gets one shared `PDF` subfolder next to those Markdown files.
- Clarified the other DEV Markdown PDF output modes: `Прямо рядом с MD`, `Общая docs/PDF`, and `Зеркало в Назначение` through the current Workbench Target.
- Replaced the old external dev-PDF folder wording in the GUI-facing documentation model.
- Compact hardware recommendation badge is now shown only in `Инструменты проекта` and `Локальный OCR`.
- Kept canonical tooltip styling and detailed control tooltips while skipping noisy per-checkbox tooltips for document-kind chips.

2026-07-01 - OCR Brick integration and UI cleanup

Added

- Added root-level OCR windows:
 - `Офисный текст без OCR`
 - `Локальный OCR`
 - `Платный API OCR`
 - `Тест качества OCR`
 - `Проверка реквизитов`
 - `Таблицы`
- Added OCR Brick pipeline under `system_core/ocr_brick` with rasterize, preprocess, cache, OCR engine dispatch and Markdown/JSON output.
- Added CPU preprocessing fallback with PIL/NumPy: JPEG denoise, autocontrast, contrast boost, unsharp mask, adaptive binarization, deskew and vertical-line removal.
- Added Real-ESRGAN ncnn-vulkan portable installer and detection from `tools/realrgan-ncnn-vulkan`.

- Added engine-aware scan cleanup profile **Очистка скана**:
 - **Авто** uses tuned defaults for the selected OCR engine.
 - **Без очистки** sends the page raw.
 - **Тяжёлый скан** uses stronger denoise, contrast and sharpening.
 - **Цифры** boosts legal/financial identifiers.
 - **Ручной** exposes the old detailed preprocessing switches in Advanced.
- Added local OCR engines:
 - Tesseract as the recommended local baseline.
 - Surya as optional slow/local OCR for cases where API cost matters more than time.
- Added paid/API OCR engines:
 - Yandex Vision OCR with selectable mode/model and OCR-safe model cache.
 - xAI vision OCR with model list refresh, curated pinned models and API smoke test.
 - Gemini and OpenAI vision OCR through the same OCR Brick vision contract.
- Added Gemini API cost/scheduling controls:
 - **Gemini tier** selector with **Standard** and **Flex** for regular OCR and one-page OCR quality tests.
 - **Gemini Batch OCR** project tool to prepare local JSONL manifests and optionally submit asynchronous Batch API jobs.
 - **Забрать Gemini Batch** project tool to poll/download finished Batch OCR results into Markdown.
- Added **2-pass Tesseract** in paid OCR:
 - For Yandex it writes a local sidecar with Tesseract text/boxes.
 - For xAI/Gemini/OpenAI it can inject Tesseract text as a fallible prompt hint.
- Added one-page OCR quality benchmark for raw/clean comparisons before large OCR runs.
- Added numeric/legal requisites checker **Проверка реквизитов** with Tesseract, Surya or vision-page locator modes and xAI verification.
- Added global tooltips for command buttons, fields, model selectors, service buttons and API-key controls.
- Added API key add/delete UI with warnings and provider-specific key selectors.
- Added pinned model storage in **config/gui_model_pins.json**.
- Added hardware recommendation badge for local OCR and project tools. It detects GPU class and recommends Vulkan, llama.cpp CUDA or PyTorch CUDA install paths.

- Added project-tools mode switcher: `Установка` / `Проверка`.

Changed

- Split the previous mixed OCR Brick window into separate local and paid OCR root windows.
- Reworked local/paid OCR dialogs into logical visual blocks: engine/model/access, preprocessing, local 2-pass, prompt/parameters, postprocessing, output and advanced.
- Replaced the visible preprocessing control field with a single engine-aware profile selector; detailed knobs are hidden until `Ручной` is selected.
- Reworked model selectors into wide dropdowns with pin/unpin/update controls and curated defaults.
- Shortened noisy operation labels and tightened the main root-menu row spacing.
- Moved coordinate-table extraction to root as `Таблицы`.
- Moved Tesseract install into `Инструменты проекта -> Установка` and marked it as recommended.
- Moved project checks, model refreshes and OCR smoke tests into `Инструменты проекта -> Проверка`.
- Updated install/build scripts for the current portable runtime layout and optional OCR engines.
- Updated DEV Markdown PDF output to support project `docs/PDF`, beside-source output and external target trees.
- Updated vision OCR prompt handling to reduce decorative Markdown and preserve source text more strictly.
- Updated Gemini OCR calls to pass `service_tier=flex` when Flex is selected and use longer timeouts for Flex queueing.
- Replaced stale user guides with current root-window documentation and moved implementation history into this changelog.
- Added `docs/GEMINI_BATCH_FLEX.md` with Gemini Standard/Flex/Batch usage, GUI mapping, artifact paths and the verified `Вопрос 239.pdf` dry-run.
- Updated GitHub README files to describe OCR Brick root windows, current local/API OCR engines, Gemini Batch/Flex and one-page quality/requisites workflow instead of the legacy Workbench-first flow.
- Added install documentation for optional OCR acceleration entries `[11]` - `[17]`, including practical RTX 50xx/40xx/Vulkan/CPU choices and the distinction between Real-ESRGAN, Surya, llama.cpp and heavy PyTorch CUDA benchmark payloads.

Removed

- Removed Paddle OCR from UI, install paths and project cleanup expectations.
- Hid/removed legacy Python OCR Workbench from the visible root navigation.
- Removed stale `config/keys` initialization from folder setup.
- Removed legacy mixed root OCR actions in favor of the new local/paid OCR windows.

Verified

- NiceGUI smoke test passes after UI changes.
- Python compile checks pass for updated UI/service files.
- Yandex OCR smoke passed with configured key/folder.
- xAI model-list and OCR smoke passed.
- Gemini Batch dry-run prepared 3 JSONL requests for `Вопрос 239.pdf` pages 6-8 without remote submission.
- Tesseract OCR Brick runs were smoke-tested on `10_Самарская_Тольятти.pdf` and `1_Псковская_Псковский.pdf`.
- Numeric check matched key Pskov identifiers with local locator plus xAI verification.

Cleanup

- Legacy implementation plans, handoff briefs and generated benchmark/log/cache artifacts were removed from the working tree after this changelog captured the relevant decisions.
- Source test package and smoke-run workspaces were removed from the portable project cleanup set.