

# Capstone Project Report & File Documentation

[!NOTE] **Dynamic Version Control:** This document is versioned under the master Git ID: \$Id: a7732e033c29e63f17009fb792feb0ab2e57a882 Lange François lanfr144@school.lu 2026/05/31 15:41:33 Lange François lanfr144@school.lu 2026/05/31 15:41:33 docs: add WSL runbook, WSL compose file, and Taiga JSON export to report catalog [PreRelease-1.0-6-ga7732e0] \$. All file versions and commit histories below are extracted directly from the live Git metadata logs.

---

## 1. Project Overview & Deliverables

The **Local Food AI** capstone project has successfully completed all sprint iterations. The system stands fully verified, containerized, and documented.

### What Has Been Done

- 1. Model Upgraded to Ollama Latest:** Transitioned from the `llama3.2:3b` model to the much more robust, large reasoning-focused **qwen2.5:7b** model (4.7 GB) with structured XML Chain-of-Thought (CoT) calculations. Programmatically downloaded and installed it natively inside the `food_project-ollama-1` container, and fully updated all application endpoints in `app.py`.
  - 2. Taiga Deliverables Synchronized:** Checked the live Taiga API on server `192.168.130.161`. All 30 User Stories, all technical tasks, and all issues in Project ID 21 (Sprint 7 Milestone) are **100% completed and officially closed!**
  - 3. Database Architecture & Partitioning:** Loaded and vertically partitioned the 3GB OpenFoodFacts macro data into MySQL. Configured matching FULLTEXT engines to search records in less than **0.04s** (averaging 90% latency reduction).
  - 4. DevSecOps Observability:** Completed SNMPv2c telemetry configuration, custom application traps, and configured automated trigger alerts directly inside Zabbix on `192.168.130.170`.
  - 5. Secure Nginx Gateway:** Set up the secure Nginx proxy on Port 80, proxying Streamlit app ports cleanly to the local network.
  - 6. Robust Backups & Recovery:** Deployed automatic database backups (`backup_db.sh`) and local offline single-node fallback capabilities (`docker-compose_skip.yml`).
  - 7. Sequential Operations Manager:** Created `manage_services.sh` to allow developers to safely stop, start, and restart all microservices in the proper dependency order without triggering redundant online ingestion sequences.
-

## 2. Project File Catalog & Documentation

Below is an exhaustive catalog of every critical file in the repository, detailing its path, functional purpose, and active Git version tags.

*Note: This chapter is compiled in landscape layout inside Project.pdf to guarantee complete columns readability.*

<b>File Path</b>	<b>Purpose &amp; Technical Responsibility</b>	<b>Commit</b>	<b>Author</b>	<b>Commit Date</b>	<b>Last Commit Message</b>
<b>app.py</b> ./app.py	Core Streamlit Web Application. Hosts the clinical food search engine, the RAG chat dietitian interface (utilizing Ollama and SearXNG tool calling), and the visual plate builder.	3e5cddb	Lange François	2026/05/22 09:47:33	<i>TG-221 #closed - Refactor Ask Chat system prompt to use Chain of Thought (CoT) reasoning for higher clinical accuracy</i>
<b>ingest_csv.py</b> ./ingest_csv.py	High-performance background database loader. Stream-reads and batch-inserts the 3GB OpenFoodFacts dataset into MySQL using Pandas chunking and optimizes indices post-load.	264d274	lanfr144	2026/05/21 09:43:09	<i>TG-442: Sync resilience configurations, resolve SearXNG crash, and update docs with dynamic custom Git log ID and tag</i>
<b>unit_converter.py</b> ./unit_converter.py	Mathematical converter engine that parses natural recipe volume inputs (e.g. cups, spoons) and converts them to metric weights based on macro density mappings.	ea04a85	lanfr144	2026/05/08 08:57:06	<i>TG-86: finalize system pre-initialization, auto-pull LLM, egg scales</i>

<b>snmp_notifier.py</b> ./snmp_notifier.py	Observability SNMP utility. Formulates and transmits raw SNMP trap payloads to the central Zabbix monitoring server on critical application failures.	264d274	lanfr144	2026/05/21 09:43:09	<i>TG-442: Sync resilience configurations, resolve SearXNG crash, and update docs with dynamic custom Git log ID and tag</i>
<b>configure_zabbix_alerts.py</b> ./configure_zabbix_alerts.py	DevOps provisioning script. Uses the Zabbix API to automatically set up host groups, custom templates, items, triggers, actions, and media types for alerts.	264d274	lanfr144	2026/05/21 09:43:09	<i>TG-442: Sync resilience configurations, resolve SearXNG crash, and update docs with dynamic custom Git log ID and tag</i>
<b>configure_zabbix_email.py</b> ./configure_zabbix_email.py	Security & Monitoring. Configures email media types and SMTP server routes for Zabbix alert notifications on system downtime.	ade82af	lanfr144	2026/05/18 14:08:27	<i>TG-196: Full security refactor, Taiga sync, and Data pipeline automation</i>
<b>zabbix_telemetry.py</b> ./zabbix_telemetry.py	Monitoring agent daemon. Queries active application statistics, memory, and query timers to supply Zabbix telemetry indicators.	ade82af	lanfr144	2026/05/18 14:08:27	<i>TG-196: Full security refactor, Taiga sync, and Data pipeline automation</i>
<b>check_users.py</b> ./check_users.py	Security utility. Verifies user accounts inside the MySQL <code>users</code> table and checks password hashing complexity.	7766898	lanfr144	2026/04/29 14:39:55	<i>Add check users script</i>
<b>rotate_passwords.py</b> ./rotate_passwords.py	Administrative credential utility. Cycles and re-encrypts database passwords within the <code>.env</code> secret file.	ade82af	lanfr144	2026/05/18 14:08:27	<i>TG-196: Full security refactor, Taiga sync, and Data pipeline automation</i>

<b>myloginpath.py</b> ./myloginpath.py	MySQL credential companion helper that simplifies the generation of encrypted login path configuration profiles.	4655c26	lanfr144	2026/04/29 08:30:03	Add untracked project files and configs
<b>data_sync.sh</b> ./data_sync.sh	Master pipeline coordinator. Supports download fetching in --online mode and local file processing in offline fallback mode.	264d274	lanfr144	2026/05/21 09:43:09	TG-442: Sync resilience configurations, resolve SearXNG crash, and update docs with dynamic custom Git log ID and tag
<b>backup_db.sh</b> ./backup_db.sh	Resiliency backup automation. Runs mysqldump on user tables inside the active container and prunes backups older than 7 days.	264d274	lanfr144	2026/05/21 09:43:09	TG-442: Sync resilience configurations, resolve SearXNG crash, and update docs with dynamic custom Git log ID and tag
<b>reset.sh</b> ./reset.sh	Teardown script. Wipes local temporary containers and prunes volume locks during crashes.	264d274	lanfr144	2026/05/21 09:43:09	TG-442: Sync resilience configurations, resolve SearXNG crash, and update docs with dynamic custom Git log ID and tag
<b>proper_reset.sh</b> ./proper_reset.sh	High-level administrative wipe script that brings the entire network stack and repositories back to a pristine state.	776d6a6	lanfr144	2026/04/29 12:44:49	Add proper reset

<b>deploy.sh</b> ./deploy.sh	Naked OS installation guide. Installs necessary system packages, Python venv libraries, and native Ollama.	a54dc25	lanfr144	2026/04/22 15:01:17	<i>TG-21: Update deploy.sh to include requests connectivity dependency.</i>
<b>start_batch_ingest.sh</b> ./start_batch_ingest.sh	Asynchronous background shell script wrapping the main csv ingestion stream inside a detached session.	00f1d63	lanfr144	2026/04/24 07:50:40	<i>Fix python virtual env paths</i>
<b>download_csv.sh</b> ./download_csv.sh	Downloader helper script that fetches specific smaller subsets of OpenFoodFacts CSV files.	1a3cdca	lanfr144	2026/05/05 07:14:54	<i>fix: resolve pip encoding issue and add exec permissions to download script</i>
<b>master_trigger.sh</b> ./master_trigger.sh	Orchestrator script that wakes and verifies multiple secondary subservices in sequence.	38a83a1	lanfr144	2026/04/23 10:50:37	<i>Deployment Finalization: Vitamin schemas, Green UI, and Taiga tools</i>
<b>manage_services.sh</b> ./manage_services.sh	DevOps service manager script. Handles automated, sequential startup, shutdown, restart, and health checking of all container elements in the stack.	78a1c2c	Lange François	2026/05/22 07:55:19	<i>docs: Hardening, hybrid landscape, documentation index, and US-203 Taiga tasks alignment</i>
<b>generate_docs.py</b> ./generate_docs.py	Dynamic doc generator. Generates and mirrors all markdown manuals under /docs with live Git log metadata injection.	09c5304	Lange François	2026/05/22 09:19:09	<i>TG-220 TG-221 TG-222 #closed - Upgrade Ollama to Qwen2.5-7B, refactor backend prompts for XML scratchpad reasoning, and implement response parsing</i>

<b>docker-compose.yml</b> ./docker-compose.yml	Main 10-container Docker orchestration map defining MySQL, App UI, Ollama Engine, SearXNG, Nginx proxy, Airflow stack, and Zabbix server suites.	264d274	lanfr144	2026/05/21 09:43:09	<i>TG-442: Sync resilience configurations, resolve SearXNG crash, and update docs with dynamic custom Git log ID and tag</i>
<b>docker-compose_skip.yml</b> ./docker-compose_skip.yml	Resilient 8-container offline/local single-node orchestration manifest.	264d274	lanfr144	2026/05/21 09:43:09	<i>TG-442: Sync resilience configurations, resolve SearXNG crash, and update docs with dynamic custom Git log ID and tag</i>
<b>docker-compose-wsl.yml</b> ./docker-compose-wsl.yml	WSL2-specific Docker Compose configuration file. Configures services with a +20 port shift to guarantee zero port conflicts on developer workstations.	c52c6a1	Lange François	2026/05/31 15:41:22	<i>docs: add WSL installation runbook, WSL compose file with shifted ports, and Taiga JSON export</i>
<b>alembic.ini</b> ./alembic.ini	Alembic configuration setting routing database connection URIs for versioning schemas.	73f7a04	lanfr144	2026/04/24 16:18:55	<i>Optimize horizontal partitioning to slice into 8-column chunks bypassing InnoDB limits</i>

<b>my.cnf</b> ./my.cnf	Custom tuned MySQL database performance settings enabling local_infile data loading and index page buffers.	86c76e2	lanfr144	2026/04/17 10:26:35	<i>TG-1: Fix MySQL 8.0 startup crash by removing premature validate_password plugin config</i>
<b>.env</b> ./env	Secret storage container holding encrypted MySQL user passwords and active environment flags.	ca3877d	lanfr144	2026/05/13 11:15:42	<i>Stop save the .env file</i>
<b>.gitattributes</b> ./gitattributes	Git clean/smudge layout mapping enabling automatic tracking of dynamic \$Id: a7732e033c29e63f17009fb792feb0ab2e57a882 Lange François lanfr144@school.lu 2026/05/31 15:41:33 Lange François lanfr144@school.lu 2026/05/31 15:41:33 docs: add WSL runbook, WSL compose file, and Taiga JSON export to report catalog [PreRelease-1.0-6-ga7732e0] \$ metadata expansion within version files.	0cfd52	lanfr144	2026/05/07 09:54:17	<i>TG-85: enable export-subst for Format string git identification</i>
<b>requirements.txt</b> ./requirements.txt	Python runtime dependency catalog storing strict library versioning constraints.	bb2ac28	lanfr144	2026/05/11 07:59:05	<i>fix requirements.txt encoding for fpdf2</i>
<b>INSTALL_WSL.md</b> ./INSTALL_WSL.md	WSL2 deployment guide. Provides step-by-step instructions for installing and deploying the application inside WSL2 with port shifts.	c52c6a1	Lange François	2026/05/31 15:41:22	<i>docs: add WSL installation runbook, WSL compose file with shifted ports, and Taiga JSON export</i>

<b>taiga/local-food-ai-1-36f35ff9-da1b-4eb5-9309-058448c998ad.json</b> ./taiga/local-food-ai-1-36f35ff9-da1b-4eb5-9309-058448c998ad.json	Historical Taiga Agile export. Contains the complete project history, including all closed user stories, tasks, and sprint configurations.	c52c6a1	Lange François	2026/05/31 15:41:22	<i>docs: add WSL installation runbook, WSL compose file with shifted ports, and Taiga JSON export</i>
<b>scripts/generate_pdfs.py</b> ./scripts/generate_pdfs.py	PDF document builder. Converts all markdown documentation manuals under /docs into high-fidelity PDF format with expanded Git version headers.	78a1c2c	Lange François	2026/05/22 07:55:19	<i>docs: Hardening, hybrid landscape, documentation index, and US-203 Taiga tasks alignment</i>
<b>scripts/generate_project_report.py</b> ./scripts/generate_project_report.py	Technical project report generator. Automatically gathers codebase structure, Git commit metadata, and purpose records to construct the Project.pdf report.	a7732e0	Lange François	2026/05/31 15:41:33	<i>docs: add WSL runbook, WSL compose file, and Taiga JSON export to report catalog</i>
<b>scripts/setup_deploy.py</b> ./scripts/setup_deploy.py	DevOps deployment script. Orchestrates local and VM container sets, verifying network connectivity and system parameters.	0065125	lanfr144	2026/05/20 08:52:08	<i>TG-202: Add log rotation limits to prevent 100% disk usage</i>
<b>scripts/taiga_sync_final.py</b> ./scripts/taiga_sync_final.py	Taiga automated synchronization helper. Pushes bug tickets, fills wiki pages, and assigns unassigned user stories.	a4342d3	lanfr144	2026/05/19 09:09:10	<i>TG-198: Add Taiga consistency automation script for full Jira/Agile alignment</i>

---

### 3. Directory Structure Map

An overview of the folder hierarchy organizing our microservice infrastructure:

- **[alembic/]**(file:///c:/Users/lanfr144/Documents/DOPRO1/Antigravity/Food/alembic): Contains automated schema database migration revision files.
  - **[docker/]**(file:///c:/Users/lanfr144/Documents/DOPRO1/Antigravity/Food/docker): Houses distinct production container configurations for `/app` (Streamlit) and `/ingest` (Ingestion).
  - **[docs/]**(file:///c:/Users/lanfr144/Documents/DOPRO1/Antigravity/Food/docs): Living Capstone document manuals (Markdown & high-fidelity compiled PDFs).
  - **[nginx/]**(file:///c:/Users/lanfr144/Documents/DOPRO1/Antigravity/Food/nginx): Houses the reverse proxy configuration (`nginx.conf`) forwarding local port 80 traffic.
  - **[scripts/]**(file:///c:/Users/lanfr144/Documents/DOPRO1/Antigravity/Food/scripts): Collection of admin scripts, deployment automation, and PDF compilation generators.
  - **[searxng/]**(file:///c:/Users/lanfr144/Documents/DOPRO1/Antigravity/Food/searxng): Core configuration files (`settings.yml`) securing private, localized search operations.
- 

### 4. Operational Next Steps (Day 2 Procedures)

1. **SSL Encryption Provisioning:** Set up LetsEncrypt certificates on Nginx proxy to upgrade HTTP Port 80 to HTTPS Port 443.
2. **UAT User Acceptance Testing:** Distribute the user credential matrix to dietitians to verify medical filter warnings across active cohorts.
3. **Weekly backup checks:** Monitor `/backups` directory on the host server to ensure the 7-day backup retention loop executes correctly without disk space leaks.