

# Local Food AI - Detailed Operator Installation Guide

Project Deliverable

---

**Course** B1AI

**Class:**

**Author:** François LANGE

**Date:** 2026/06/19



# Table of Contents

- • Pre-Deployment Operator Survey (Pre-requisites Gathering) ..... Page 3
- • Platform Mapping: Which Container Goes Where? ..... Page 3
- • Platform Provisioning Commands ..... Page 4
  - 3.1: WSL2 Provisioning (Local Client Workstation) ..... Page 4
- Run in Administrator PowerShell ..... Page 4
  - 3.2: Hyper-V VM Provisioning (Server A - Database & Zabbix) ..... Page 4
- Run in Administrator PowerShell on Server A ..... Page 4
  - 3.3: VirtualBox VM Provisioning (Server B - Ollama AI Engine) ..... Page 4
- Run in Command Prompt on Server B ..... Page 4
  - 3.4: Native Physical Unix/Linux Server (No Virtualization) ..... Page 4
    - • Secure Authentication & SSH Exchange ..... Page 5
- • Generate SSH Keys on WSL Client ..... Page 5
- • Push Key to Database VM (Server A) ..... Page 5
- • Push Key to AI VM (Server B) ..... Page 5
  - • Multi-Node Docker Network & Configuration ..... Page 2
    - Step 5.1: Configure WSL Client `.env` ..... Page 2
    - Step 5.2: Configure Ollama (VirtualBox Server B) Listening Port ..... Page 2
- SSH into Server B (192.168.130.161) ..... Page 5
- Add the environment variables: ..... Page 5
- Reload and restart service ..... Page 6
  - • Zabbix Reconfiguration for Multi-Node SNMPv3 Telemetry ..... Page 2

- Step 6.1: Deploy SNMPv3 Daemons ..... Page 2
- Listen on all interfaces ..... Page 2
- Create secure SNMPv3 User ..... Page 2
  - Step 6.2: Configure Zabbix Server Dashboard (Web UI) ..... Page 2
    - Verifying Alert Channels ..... Page 7
    - 7.1: Microsoft Teams / Discord Alert Webhook ..... Page 2
    - 7.2: Password Reset Email (SMTP Gateway) ..... Page 2
  - Operator Post-Installation Checklist ..... Page 7

```

The current version is #ident
"@(#)$Format:LocalFoodAI_lanfr144:generate_docs.py:Francois
Lange:lanfr144@school.lu:2026/06/16 21:48:22:Francois
Lange:lanfr144@school.lu:2026/06/16
21:48:22:2a8ed056889f3b796f9266feda591b12b72f3b96:HEAD -> main, origin/main:$"
  
```

# Local Food AI - Detailed Operator Installation Guide

This document is a step-by-step installation, configuration, and verification manual for deploying the **Local Food AI** system in an enterprise environment. It covers mixed environments ranging from native, physical Unix/Linux servers (bare-metal) to hybrid hypervisor infrastructures (WSL2, Hyper-V, and VirtualBox), cross-node networking, SNMPv3 monitoring, alert channels, and acceptance testing.

## 1. Pre-Deployment Operator Survey (Pre-requisites Gathering)

Before running installation scripts, the operator **must** collect the following physical/virtual infrastructure parameters and store them in the deployment matrix:

REQUIRED PARAMETER	OPERATOR INPUT / DESCRIPTION
Deployment Workstation IP	e.g., 192.168.1.50
Hyper-V Host VM IP	e.g., 192.168.130.170
VirtualBox Host VM IP	e.g., 192.168.130.161
SSH Key Location (Private)	e.g., ~/.ssh/id_rsa
SMTP Relay Password	e.g., ***** (For Zabbix/App password reset email)
Teams/Discord Webhook URL	e.g., https://discord.com/api/webhooks/...

## 2. Platform Mapping: Which Container Goes Where?

To maximize CPU/GPU efficiency and secure database read/writes, services can be distributed across native physical servers or virtualized environment hypervisors:

COMPONENT CONTAINER	DEPLOYMENT ENVIRONMENT	WHY
streamlit-app (app.py)	Local WSL2 or Bare-Metal Unix	Low-latency rendering and direct client access
mysql (Database Node)	Bare-Metal Server or Hyper-V VM	Persistent enterprise-grade disk storage
ollama (NLP llama3.2:3b Engine)	Bare-Metal Server or VirtualBox VM	Dedicated CPU/GPU resource allocation

<b>zabbix-server &amp; web (Monitoring)</b>	Bare-Metal Server or Hyper-V VM	Centralized SNMPv3 alert processing and logs
<b>searxng (Meta-Search Gateway)</b>	Local WSL2 or Bare-Metal Unix	Dynamic browser-level loopbacks

## 3. Platform Provisioning Commands

### 3.1: WSL2 Provisioning (Local Client Workstation)

Enable WSL2 and install Ubuntu 24.04:

```
# Run in Administrator PowerShell
dism.exe /online /enable-feature /featurename:Microsoft-Windows-Subsystem-Linux /all /norestart
dism.exe /online /enable-feature /featurename:VirtualMachinePlatform /all /norestart
wsl --install -d Ubuntu-24.04
```

### 3.2: Hyper-V VM Provisioning (Server A - Database & Zabbix)

Deploy a dedicated Ubuntu VM on Hyper-V using PowerShell:

```
# Run in Administrator PowerShell on Server A
New-VM -Name "FoodAI-Database-Node" -MemoryStartupBytes 8GB -Generation 2 -NewVHDPATH "C:\VMs\FoodAI_DB.vhdx" -VHDSIZEBYTES 80GB -SwitchName "External Switch"
Set-VMFirmware -VMName "FoodAI-Database-Node" -EnableSecureBoot Off
Start-VM -Name "FoodAI-Database-Node"
```

### 3.3: VirtualBox VM Provisioning (Server B - Ollama AI Engine)

Deploy a dedicated VM on VirtualBox using Command Line:

```
# Run in Command Prompt on Server B
vboxmanage createvm --name "FoodAI-AI-Node" --ostype "Ubuntu_64" --register
vboxmanage modifyvm "FoodAI-AI-Node" --memory 8192 --cpus 4 --vram 128 --nic1 bridged --bridgeadapter1 "Intel Ethernet Connection"
vboxmanage createhd --filename "C:\VMs\FoodAI_AI.vdi" --size 60000
vboxmanage storagectl "FoodAI-AI-Node" --name "SATA Controller" --add sata --controller IntelAHCI
vboxmanage storageattach "FoodAI-AI-Node" --storagectl "SATA Controller" --port 0 --device 0 --type hdd --medium "C:\VMs\FoodAI_AI.vdi"
vboxmanage startvm "FoodAI-AI-Node" --type headless
```

### 3.4: Native Physical Unix/Linux Server (No Virtualization)

For a native physical installation directly on the hardware (bare-metal Ubuntu Server):

- Install Ubuntu Server 22.04 LTS or 24.04 LTS on the physical machine.
- Ensure SSH is installed and running:

```
sudo apt update && sudo apt install -y openssh-server
```

```
sudo systemctl enable --now ssh
```

- Configure a static IP address and ensure the host firewall allows traffic on the required ports (e.g. 80, 8502, 3306, 8081).

## 4. Secure Authentication & SSH Exchange

Exchange SSH public keys to allow automated, passwordless container management across nodes:

```
bash
# 1. Generate SSH Keys on WSL Client
ssh-keygen -t rsa -b 4096 -f ~/.ssh/id_rsa_foodai -N ""

# 2. Push Key to Database VM (Server A)
ssh-copy-id -i ~/.ssh/id_rsa_foodai.pub operator@192.168.130.170

# 3. Push Key to AI VM (Server B)
ssh-copy-id -i ~/.ssh/id_rsa_foodai.pub operator@192.168.130.161
```

## 5. Multi-Node Docker Network & Configuration

To allow WSL, Hyper-V, and VirtualBox nodes to communicate, update the `.env` variables and `docker-compose.yml` to use bridged network endpoints.

### Step 5.1: Configure WSL Client `.env`

Update `.env` in the Streamlit workspace:

```
ini
DB_HOST=192.168.130.170
DB_USER=food_reader
DB_PASS=reader_pass
APP_AUTH_USER=food_app_auth
APP_AUTH_PASS=auth_pass
OLLAMA_HOST=http://192.168.130.161:11434
SEARXNG_HOST=http://localhost:8080
ZBX_SERVER_HOST=192.168.130.170
```

### Step 5.2: Configure Ollama (VirtualBox Server B) Listening Port

Ensure the Ollama daemon inside VirtualBox binds to `0.0.0.0` (all interfaces):

```
bash
# SSH into Server B (192.168.130.161)
sudo systemctl edit ollama.service

# Add the environment variables:
```

```
[Service]
Environment="OLLAMA_HOST=0.0.0.0"

# Reload and restart service
sudo systemctl daemon-reload
sudo systemctl restart ollama
```

## 6. Zabbix Reconfiguration for Multi-Node SNMPv3 Telemetry

To monitor all distributed deployment environments securely:

### Step 6.1: Deploy SNMPv3 Daemons

Install and configure SNMPv3 daemons on WSL, Hyper-V Database VM, and VirtualBox AI VM:

```
bash
sudo apt update && sudo apt install -y snmpd
```

Edit `/etc/snmp/snmpd.conf`:

```
# Listen on all interfaces
agentAddress udp:161

# Create secure SNMPv3 User
createUser securityUser SHA "securityAuthPassword" AES "securityPrivPassword"
rouser securityUser authpriv
```

Restart daemon:

```
bash
sudo systemctl restart snmpd
```

### Step 6.2: Configure Zabbix Server Dashboard (Web UI)

- Open Zabbix in your browser at <http://192.168.130.170:8081>.
- Navigate to **Configuration > Hosts > Create Host**.
- Create three distinct hosts:
  - **WSL-Workstation** (IP: [192.168.1.50](#))
  - **Database-Node** (IP: [192.168.130.170](#))
  - **AI-Node** (IP: [192.168.130.161](#))
- Add the **SNMP Interface** pointing to Port 161 for each host.
- In the **Security Tab**, select SNMPv3, enter Username `securityUser`, select Auth Protocol `SHA` / `securityAuthPassword`, and Privacy Protocol `AES` / `securityPrivPassword`.

- Attach the pre-installed **Local Food AI Telemetry** Template.

## 7. Verifying Alert Channels

### 7.1: Microsoft Teams / Discord Alert Webhook

To verify Zabbix is communicating with Discord / Teams:

- Trigger a test CPU threshold spike inside WSL:

```
yes > /dev/null & sleep 10 ; killall yes
```

- Verify Zabbix triggers the alert and transmits the notification.
- Check your designated channel for the incoming payload:
  - Expected Output: `[PROBLEM] High CPU Utilization Detected on WSL-Workstation.`

### 7.2: Password Reset Email (SMTP Gateway)

- In the Streamlit UI Sidebar, select **Reset Password**.
- Trigger a reset link for user `ClinicianA`.
- Check the inbox or SMTP system log (`tail -f /var/log/mail.log` on Server A) to verify outbound delivery.

## 8. Operator Post-Installation Checklist

Run these test cases to verify the installation:

TEST CASE ID	ACTIONS TO PERFORM	EXPECTED RESULTS	STATUS
TC-OP-01	Search 'Cheese' on Search Tab	10+ records returned in <0.04s. Listeria warning flags on unpasteurized.	[ ]
TC-OP-02	Enter '1.5 cups' in Plate Tab	Parsed and converted to metric grams based on density index.	[ ]
TC-OP-03	Ask Chat: 'Can I eat sushi?'	llama3.2:3b retrieves database context and flags raw fish as forbidden for pregnancy.	[ ]
TC-OP-04	Trigger manual db backup	Timestamped compressed .sql.gz created inside backups/ folder.	[ ]
TC-OP-05	Terminate Ollama Container	Zabbix PROBLEM active alert generated on dashboard in < 30 seconds.	[ ]

## References

- **OpenFoodFacts Dataset & API Catalog**: Detailed food ingredients database. (<https://>

[world.openfoodfacts.org/](https://world.openfoodfacts.org/))

- **Ollama Local LLM Inference Engine:** Lightweight instruction-following llama3.2 runtimes. (<https://ollama.com/>)
- **Zabbix Enterprise Telemetry and Monitoring:** System health and performance logging. (<https://www.zabbix.com/>)

## Index

- **AI:** Page 1, Page 2, Page 3, Page 4, Page 5, Page 6, Page 7, Page 9
- **MySQL:** Page 3, Page 9
- **Zabbix:** Page 1, Page 2, Page 3, Page 4, Page 6, Page 7, Page 8, Page 9
- **Docker:** Page 1, Page 5, Page 9
- **Streamlit:** Page 3, Page 5, Page 7, Page 9
- **Nginx:** Page 9
- **RAG:** Page 3, Page 4, Page 9
- **Allergens:** Page 9
- **Vitamins:** Page 9
- **Minerals:** Page 9
- **Clinical:** Page 9
- **WSL:** Page 1, Page 3, Page 4, Page 5, Page 6, Page 7, Page 9
- **Ollama:** Page 1, Page 3, Page 4, Page 5, Page 6, Page 7, Page 8, Page 9
- **LLM:** Page 8, Page 9
- **Database:** Page 1, Page 3, Page 4, Page 5, Page 6, Page 7, Page 9
- **Security:** Page 6, Page 9
- **Telemetry:** Page 1, Page 6, Page 7, Page 8, Page 9
- **Backup:** Page 7, Page 9
- **Firewall:** Page 9
- **SMTP:** Page 2, Page 3, Page 7, Page 9