Edge-Interoperable AI & ML Blueprint

Bringing Intelligent Services to Africa's Communities — At the Edge

LOUIS FENDJI, UNIVERSITY OF NGAOUNDÉRÉ, CAMEROON

Funded by the European Union

DIGITAfrica project has received funding from the EU Horizon Europe research and innovation Programme and Switzerland under Grant Agreement No. 101187966. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Executive Agency (REA). Neither the European Union nor the granting authority can be held responsible for them.

DIGIT AFRICA

The AI Opportunity — and Africa's Specific Challenge

Why AI must be brought to the edge, not kept in the cloud

Cloud Dependency

Most AI services rely on centralised cloud infrastructure abroad — creating latency, cost, and sovereignty risks for African users

Connectivity Gaps

Unreliable or absent internet at clinics, farms, and schools means cloud-based AI simply cannot serve where it's needed most

Data Sovereignty

Sensitive health, agricultural, and education data must not leave the country or institution — local AI processing is essential

Skills & Ownership

African researchers and institutions must build, train, and own their AI models — not just consume products built elsewhere

What is the Edge AI & ML Blueprint?

A reusable, open-source framework for sovereign AI at the edge

Edge AI Blueprint

A framework for deploying reproducible AI and ML workflows across distributed edge-cloud infrastructures — enabling AI services to run close to the data they serve.

Built on composable baseline services. Deployable on a single node. Designed for data sovereignty and autonomous operation without cloud dependency.

Free • Open Source • Africa-Led

Reproducible
ML workflows run identically across all partner sites — same results whether in Tunis or Ngaoundéré.

Sovereign
Data never leaves the institution. Models are trained and served locally. No external dependency.

Interoperable
Shares a common service pool with the Heterogeneous Networking Blueprint. Federation-ready from day one.

What the Edge AI Blueprint Does

From data collection to model deployment — entirely at the edge

Run complete AI/ML workflows — data ingestion, model training, inference, and model sharing — on low-cost edge hardware, without sending data to the cloud.

Core Capabilities

- Local AI training & inference on GPU/accelerator hardware
- Federated learning across distributed partner sites
- MLflow model lifecycle management & sharing
- JupyterHub multi-user notebook environment
- FAIR-compliant data storage (MinIO, NFS, CEPH)
- Kubernetes/K3s orchestration for workload management

Why It Matters for Africa

Privacy	Sensitive health and agricultural data stays on-site — no cloud upload required
Latency	Local inference means real-time responses at rural clinics and farms, even offline
Cost	Open-source stack on commodity hardware — no cloud compute bills or licence fees
Capacity	Researchers and students build and own their models, growing local AI expertise

Progressive Deployment: From Rural School to Pan-African Federation

Start small. Grow as capacity allows.

Tier 0

Rural School or Clinic

Offline AI inference
Local training
AI accelerator
Solar-powered

$1–1.5k

Tier 1

Campus Lab

Edge AI & inference
Multi-user access
Integrates Tier-0 nodes
AI accelerator

$2–5k

Tier 2

Regional Hub

Distributed AI workloads
Model lifecycle mgmt
Redundant backhaul

$15–30k

Tier 3

National

Large-scale AI/ML
GPU cluster
CI/CD pipelines
NREN high-speed links

$100k+

Tier 4

Pan-African Federation

Cross-NREN
Federated models
Shared datasets
Federated identity

Shared

What the Blueprint Enables

Real AI use cases across five DIGITAfrica countries

ZA South Africa	Edge telemedicine & federated diagnostics — AI models at rural clinics, no data leaving the site
SN Senegal	AI education labs & CI/CD workflows at UCAD — students training and deploying their own models
KE Kenya	Edge AI for precision agriculture — crop disease detection and yield prediction at the farm
CM Cameroon	Localised AI research & innovation at Ngaoundéré — low-cost GPU-accelerated research infrastructure
TN Tunisia	Edge-to-cloud AI integration for health and IoT at UMA — connecting local deployments to NREN

Each use case is grounded in real community needs — health, food security, education etc.

Under the Hood: Blueprint Architecture

Four composable layers — deployable on a single node or a full cluster

USER ACCESS & ENGAGEMENT
JupyterHub · Keycloak OIDC · SLICES-RI Federation · Multi-community support

EDGE AI CAPABILITIES
Local training & inference · Federated learning · GPU/accelerator support · MLflow

STORAGE & DATA INTEGRATION
MinIO / NFS / CEPH · FAIR-compliant · Object storage for model sharing across sites

EDGE ORCHESTRATION
Kubernetes / K3s · Heterogeneous compute · Workload placement · Docker Compose

PHYSICAL INFRASTRUCTURE
Commodity x86 · Raspberry Pi 4 · GPU/AI accelerator cards · Bare-metal or VM

Design Principles

- Shared service pool with Networking Blueprint
- Composable: chain services to deploy full blueprint
- Federated learning preserves data locality
- MLflow tracks experiments & shares models
- Works fully offline — no cloud required

Current Status: Already Deployed and Running

Proof-of-concept validated at multiple partner sites

5
Partner Countries with Active Deployments

100%
Kubernetes / K3s Deployment Complete

100%
OIDC Authentication Integrated

Tier 0 + Tier 1
Validated in Multiple Environments

Component	Status	Notes
Kubernetes / K3s cluster	✔ 100%	Automated bare-metal and cluster-based deployment
JupyterHub for ML notebooks	✔ 100%	Live at jupyter.thebluprintfactory.org & digitafricasn.ucad.sn
OIDC authentication (Keycloak)	✔ 100%	Integrated with SLICES-RI federation
NFS & local storage	🔄 80%	Long-term cross-site sync in progress
MLflow model sharing via S3	⏳ Dev	Model-sharing across sites under active development
Tier 2 / 3 (regional & national scale)	⏳ Dev	Planned for WP3 Phase 2 (Month 18–24)
Smartphone / mobile Tier-1 access	⏳ Dev	Expanding Tier-1 to mobile devices

How to Get Involved

Three pathways — pick the one that matches your organisation's readiness

Early Adopter

Deploy Blueprint Tier 0 or Tier 1

- Access the open GitLab repository today
- Run AI inference on a single laptop or RPi
- Join the DIGITAfrica community of practice
- Receive technical support from WP2 team

Research Partner

Co-develop & validate AI use cases

- Contribute a health, agriculture, or education AI use case
- Co-author publications and deliverables
- Shape the federated learning roadmap
- Access SLICES-RI experimental AI infrastructure

Strategic Partner

Champion, or mandate adoption

- Support national AI blueprint deployment
- Integrate blueprints into AI policy frameworks
- Provide GPU hardware, spectrum, or compute
- Anchor pan-African AI federation at Tier 4

Africa's AI future should be built in Africa. Let's start now.

Explore the blueprint: gitlab.inria.fr/digitafrica/blueprints

Project site: digitafrica.eu

Join the hands-on session tomorrow

On behalf of the DIGITAfrica consortium

DIGIT AFRICA

SORBONNE UNIVERSITÉ

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Consiglio Nazionale delle Ricerche

Inria

University of Thessaly
Founded 1984

Strathmore University

Technische Universität Berlin

Université Cheikh Anta Diop de Dakar

جامعة منوبة
Université de la Manouba

Université de Ngaoundéré
The University of Ngaoundéré

University of Amsterdam

University of Cape Town
Universiteit van Kaapstad
Yunivesithi yaseKapa

MANDAT INTERNATIONAL

Thank you!