



# Workforce Lite

## Build and secure on-prem AI for mission-critical retrieval-augmented generation (RAG)

Workforce Lite is an open-source AI-agent platform that ships DISA-STIG-hardened editions of Ollama (LLM runner), OpenSearch (vector/RAG index), and OpenWebUI (chat/notebook UI) in a single container bundle. Teams stand it up with Docker Compose or Helm, load any quantized or custom-trained model, and spin up Retrieval-Augmented Generation over mission documents—without ever leaving the enclave or paying a license fee.

### Mission Alignment



#### Mission Need

Analysts require GenAI teammates that reside 100% inside classified or restricted networks.



#### The Challenge

Commercial chatbots force data upload, lack BYO-model support, and fail DoD accreditation.



#### Our Solution

Workforce Lite installs in minutes as a self-hosted, guard-railed RAG workspace that keeps models, data, and outputs under government control.

## Key Benefits



Stand-up in minutes via one Docker Compose file



Run any quantized or GPU-accelerated LLM on-prem via GPU or CPU



Enable RAG over mission documents with one click



Enforce RBAC & LDAP for team sharing



Containers SCAP-scanned to speed ATO

## About Agile Defense

Agile Defense stands at the forefront of innovation, driving advanced capabilities and solutions tailored to the most critical national security and civilian missions. Whether developing specialized solutions, contextualizing data, or strengthening cybersecurity, our expertise is instrumental in safeguarding our nation's sensitive assets.

## Feature Breakdown

### BYO-Model & Data Engine

Ollama manages local LLMs while OpenSearch indexes embeddings for high-speed RAG. Administrators pull open-weights or load fine-tuned models with a single command, then ingest PDFs, Office docs, or API feeds through built-in pipelines—creating mission-specific knowledge bases with no external calls.

### Secure, Collaborative Web Interface

OpenWebUI offers chat, notebook, and agent-orchestration views that run entirely offline. Role-based access control integrates with existing LDAP, letting teams share agents, models, and knowledge bases while audit logs capture every interaction for oversight and continuous compliance.

### Hardened Zero-Trust Footprint

Each container is STIG-aligned, CVE-scanned, and signed, simplifying RMF and Continuous ATO packages. The bundle runs identically on Kubernetes, Docker Compose, or a standalone laptop, ensuring the same protected experience in cloud regions, SCIFs, or forward-deployed edge nodes.

### Bring It On →

Connect & Engage with Agile Defense!

### Agile Defense

1430 Spring Hill Road  
Suite #200  
McLean, VA 22102

### Contact Us

(703) 351-9977  
agiledefense.com  
info@agile-defense.com

### Learn More

Bring your unique experiences to our collective potential →

