NULLVECTOR
.monster

PHASE 07 — DEPLOYMENT

DAYS 127–150 · MLOps · PORTFOLIO · AZURE AI-102 · 10 APPLICATIONS

The final phase. Everything you built becomes something an employer can verify.

Phase 7 is not about learning new skills — it is about proving the skills you have. You will add production-grade MLOps to your best projects. You will build a portfolio that tells a clear story. You will earn the Azure AI Engineer Associate certification. And on Day 150, you will send 10 job applications with a portfolio that makes hiring managers stop and read.

Day 150 is the finish line. It is also the starting line. The 150 days built the foundation. What you build next is up to you. The difference between Day 150 and Day 0 is not just skills — it is the proven ability to show up every day and build something.

// PHASE 7 AT A GLANCE

Duration	24 days (Days 127–150) · 3.5 hours per day
Milestone	Day 147: Azure AI-102 cert. Day 150: 10 applications sent.
Azure cert	learn.microsoft.com/en-us/credentials/certifications/azure-ai-engineer/
MLOps	mlops.community — start here
Docker	docs.docker.com/get-started/

CI/CD	**GitHub Actions — docs.github.com/en/actions**
Job research	**levels.fyi/t/machine-learning-engineer**
Your library	**Essential Guide to LLMOps + LLM Design Patterns (your Humble Bundle)**

WEEK 1 · DAYS 127–133 · MLOPS — PRODUCTION AI

DAY 127 · MLOPS FUNDAMENTALS

How AI systems work after you close your laptop.

■ INDUSTRY (15 MIN)

Read: MLOps Community — browse recent articles. Note what practitioners are talking about.

→ mlops.community

■ STUDY (90 MIN) — MLOps overview

MLOps = ML + DevOps. The discipline of deploying and maintaining ML systems reliably.

Read the MLOps Maturity Model: assess where your current projects sit (level 0–3).

Key capabilities: reproducibility, automation, monitoring, retraining pipelines.

Read: 'What is MLOps' — ml-ops.org

■ BUILD (90 MIN) — Reproducible training environment

Pick your best Phase 3 or Phase 4 model. Make the training fully reproducible:

requirements.txt with pinned versions (pip freeze > requirements.txt)

config.yaml with all hyperparameters (no magic numbers in code)

Set random seeds everywhere (torch.manual_seed, numpy.random.seed)

README with exact commands to reproduce training

Verify: delete everything, follow your own README, get the same result.

✓ DONE WHEN: Training is fully reproducible from README alone. requirements.txt and config.yaml committed.

DAY 128 · DOCKER — CONTAINERIZE YOUR MODELS

Runs on your machine AND everywhere else.

■ INDUSTRY (15 MIN)

Search: 'why Docker changed software deployment explained simply'

Understanding why Docker was invented helps you understand what it solves.

■ STUDY (90 MIN) — Docker for ML

Container: a package that includes your code AND its entire environment.

Dockerfile: the recipe for building a container.

Key commands: docker build, docker run, docker push, docker pull.

Why ML needs Docker: 'it works on my machine' is not production.

Install Docker Desktop: docker.com/products/docker-desktop

→ docs.docker.com/get-started/

■ BUILD (90 MIN) — Dockerize your RAG app

Write a Dockerfile for your Phase 5 RAG application:

FROM python:3.11-slim

WORKDIR /app

COPY requirements.txt .

RUN pip install -r requirements.txt

COPY . .

CMD ['python', 'app.py']

Build the image: docker build -t my-rag-app .

Run it: docker run -p 7860:7860 my-rag-app

Verify: app works exactly the same inside the container.

✓ DONE WHEN: RAG app containerized. Docker image builds and runs. App works identically inside container.

DAY 129 · CI/CD FOR ML

Automate the path from code to deployment.

■ INDUSTRY (15 MIN)

Search: 'CI/CD pipeline for machine learning explained simply'

■ STUDY (90 MIN) — GitHub Actions for ML

CI: Continuous Integration — automatically test every code push.

CD: Continuous Deployment — automatically deploy when tests pass.

GitHub Actions: free CI/CD built into GitHub. Defined in .github/workflows/

For ML: run tests → check model performance → deploy if both pass.

Work through the GitHub Actions quickstart.

→ docs.github.com/en/actions/quickstart

■ BUILD (90 MIN) — Automated test and deploy pipeline

Create .github/workflows/ml-pipeline.yml for your RAG app:

Trigger: on every push to main branch

Job 1: run Python tests (pytest)

Job 2: if tests pass, build Docker image

Job 3: if image builds, deploy to HuggingFace Spaces

Notification: GitHub sends you an email when it fails

Push a change. Watch the pipeline run automatically.

✓ DONE WHEN: CI/CD pipeline running. Push to main triggers automatic test and deploy.

DAY 130 · MODEL MONITORING

Know when your model starts failing.

■ INDUSTRY (15 MIN)

Search: 'data drift model drift ML monitoring why it matters'

Models degrade over time as the world changes. Monitoring catches it.

■ STUDY (90 MIN) — ML monitoring

Data drift: the distribution of inputs changes over time.

Model drift: model performance degrades because the world changed.

Infrastructure monitoring: latency, error rate, throughput.

Model monitoring: prediction distribution, confidence scores, user feedback.

Tools: MLflow, Weights & Biases, Evidently AI (pip install evidently).

■ BUILD (90 MIN) — Monitoring dashboard

Add monitoring to your deployed RAG application:

Log every request: timestamp, query, response time, response length

Log user feedback if you have a thumbs up/down button

Write a script that reads the logs and generates a daily summary

Add an alert: if error rate > 5% in 1 hour, print an alert

This is simplified. Real systems use Datadog or Grafana.

✓ DONE WHEN: Monitoring logs running. Daily summary script works. Error rate alerting implemented.

DAY 131 · CLOUD DEPLOYMENT — AWS AND GCP

Deploy beyond HuggingFace.

■ INDUSTRY (15 MIN)

Search: 'AWS vs GCP vs Azure for machine learning 2025 comparison'

■ STUDY (90 MIN) — Cloud deployment options

AWS: SageMaker endpoints (powerful, complex, expensive at scale).

AWS Lambda: serverless functions — great for lightweight inference.

GCP Cloud Run: deploy any container, scale to zero when idle — cost-effective.

GCP Vertex AI: managed ML platform similar to SageMaker.

For small models: GCP Cloud Run is the easiest and cheapest.

■ BUILD (90 MIN) — Deploy to GCP Cloud Run

Deploy your containerized model to Google Cloud Run:

Create GCP account (free tier available)

Install gcloud CLI

gcloud run deploy my-model --source . --region us-central1

Test the live endpoint with curl

You now have a model running on Google's infrastructure.

Delete when done to avoid charges.

✓ DONE WHEN: Model deployed to GCP Cloud Run. Live endpoint tested with curl. Cost noted and service deleted.

DAY 132 · AZURE AI-102 CERT PREP BEGINS

Your fifth certification.

■ INDUSTRY (15 MIN)

Search: 'Azure AI Engineer what they do and what they earn'

■ STUDY (2 HRS) — Azure AI-102 overview

Download the AI-102 exam objectives. Study all 5 skill domains:

1. Plan and manage an Azure AI solution (15–20%)

2. Implement content moderation solutions (10–15%)

3. Implement computer vision solutions (15–20%)

4. Implement NLP solutions (30–35%)

5. Implement knowledge mining and document intelligence (10–15%)

Your Humble Bundle includes an Azure AI-102 course — start it today.

→ learn.microsoft.com/en-us/credentials/certifications/azure-ai-engineer/

■ BUILD (45 MIN) — Azure AI services hands-on

Using the Azure free tier:

Create an Azure Cognitive Services resource

Use the Language service to analyze sentiment on 10 different texts

Use the Computer Vision service to describe 5 images

Use the Translator to translate 5 phrases to 3 languages each

Hands-on experience makes exam questions click.

✓ DONE WHEN: AI-102 study started. Azure Cognitive Services explored hands-on.

DAY 133 · MLOPS CAPSTONE

Production-grade from end to end.

■ INDUSTRY (15 MIN)

Search: 'MLOps best practices from Netflix Spotify Uber engineering blogs'

■ STUDY (30 MIN)

Review your Days 127–132 builds. What is missing from a production system?

The answer is usually: better tests, better monitoring, better documentation.

■ BUILD (2.5 HRS) — Production pipeline

Build a complete production ML pipeline for your best model:

Containerized with Docker

CI/CD via GitHub Actions

Monitoring with request logging and daily summary

Config management via config.yaml

Comprehensive README with deployment instructions

Cost tracking: estimated cost per 1000 requests

This is Portfolio Project #10 — your most production-ready work.

✓ DONE WHEN: Complete production pipeline deployed. All components working. Portfolio Project #10 committed.

WEEK 2 · DAYS 134–140 · PORTFOLIO + AZURE AI-102

DAY 134 · PORTFOLIO STRATEGY

Tell a story, not a list.

■ INDUSTRY (15 MIN)

Search: 'how AI hiring managers evaluate portfolios' — hear it from them directly.

■ STUDY (60 MIN) — Portfolio architecture

Your portfolio tells a story: I started from zero. Here is what I built. Here is what I can do.

Choose your 3 best projects. Be ruthless. Quality over quantity.

Each project needs: one-sentence description, the problem it solves,

your role, the technology used, measurable results, a live link.

The live link is non-negotiable. If they can't click it, it doesn't exist.

■ BUILD (2 HRS) — Choose and document top 3 projects

Select your 3 best projects from Phases 2–7.

For each, write:

Problem: one sentence on what problem it solves

Approach: two sentences on how you built it

Results: one measurable outcome (accuracy %, user count, latency, etc.)

Link: the live URL

These 3 paragraphs are the core of your portfolio.

✓ DONE WHEN: 3 best projects selected. Problem/approach/results written for each. Live links verified.

DAY 135 · TECHNICAL CASE STUDIES

Show your thinking, not just your output.

■ INDUSTRY (15 MIN)

Read one Andrej Karpathy blog post. Notice: he shows his reasoning, not just results.

→ karpathy.github.io

■ STUDY (30 MIN)

Read 2 ML engineering blog posts on Medium or Towards Data Science.

Note the structure: problem → approach → what failed → what worked → results → next steps.

■ BUILD (2.5 HRS) — Write 2 technical case studies

Write a 500-word case study for 2 of your 3 portfolio projects:

The problem and why it matters

What you tried first (and why it didn't work well enough)

What you changed and why

Final results with numbers

What you would do differently with more time

Commit as CASE_STUDY_[project].md in each project's repo.

Hiring managers read these. Make them worth reading.

✓ DONE WHEN: 2 technical case studies written and committed. Each 500 words. Results include numbers.

DAY 136 · GITHUB PROFILE OVERHAUL

Your public face as a developer.

■ INDUSTRY (15 MIN)

Search: 'how to optimize your GitHub profile for ML jobs 2025'

■ STUDY (30 MIN)

Look at 3 strong ML engineering GitHub profiles.

What do they have: pinned repos, profile README, green commits, clear descriptions.

■ BUILD (2.5 HRS) — Full GitHub overhaul

Profile README (create: [username]/[username]/README.md):

Who you are (1 sentence)

What you build (2 sentences)

Your 3 portfolio projects with links

All certifications with credential IDs

Contact information

Pin your 3 best repos. Each pinned repo needs:

Clear description, README with demo GIF or screenshot, live link

Goal: 150 green squares visible. Each project repo polished.

✓ DONE WHEN: GitHub profile README live. 3 repos pinned and polished. All certifications listed.

DAY 137 · AZURE AI-102 DEEP STUDY — NLP DOMAIN

Azure AI-102 certification preparation.

■ INDUSTRY (15 MIN)

Search: 'azure ai-102 deep study — nlp domain azure AI services 2025'

■ STUDY (90 MIN)

The NLP domain is 30–35% of the exam — the largest single domain.

Azure Language service: sentiment analysis, entity recognition, key phrases, language detection.

Azure Conversational Language Understanding (CLU): build intent classifiers and entity extractors.

Question Answering: build FAQ bots from documents.

Custom Text Classification: train Azure's models on your own labels.

Do all Microsoft Learn modules for the NLP domain.

■ BUILD (60 MIN)

Build a complete NLP pipeline using Azure Language service:

Analyze sentiment for 20 customer reviews

Extract key phrases from each

Identify named entities

Classify intent (question, complaint, compliment, other)

Output a structured summary: average sentiment, top 5 key phrases, entity types found.

✓ DONE WHEN: Azure NLP pipeline built. All 4 services used. Structured summary output.

DAY 138 · AZURE AI-102 DEEP STUDY — COMPUTER VISION DOMAIN

Azure AI-102 certification preparation.

■ INDUSTRY (15 MIN)

Search: 'azure ai-102 deep study — computer vision domain azure AI services 2025'

■ STUDY (90 MIN)

Computer Vision domain: 15–20% of exam.

Azure Computer Vision: image analysis, OCR, spatial analysis.

Azure Custom Vision: train your own image classifiers and object detectors.

Azure Face API: face detection, verification, identification.

Document Intelligence: extract structured data from forms, invoices, receipts.

Do all Microsoft Learn modules for the computer vision domain.

■ BUILD (60 MIN)

Build a document processing pipeline using Azure Document Intelligence:

Upload 5 different documents (PDF forms, invoices, or any structured document)

Use Document Intelligence to extract key fields

Output a structured JSON for each document

Compare accuracy against manual extraction

✓ DONE WHEN: Document processing pipeline extracting structured data from 5 real documents.

DAY 139 · AZURE AI-102 — PRACTICE EXAM

Azure AI-102 certification preparation.

■ INDUSTRY (15 MIN)

Search: 'azure ai-102 — practice exam azure AI services 2025'

■ STUDY (90 MIN)

Full review of all 5 exam domains.

Key Azure AI services to know by heart:

Cognitive Services: umbrella for all Azure AI APIs

Azure OpenAI: access to GPT-4, DALL-E, Whisper via Azure

Azure Machine Learning: end-to-end ML platform

Bot Framework: build conversational AI

Responsible AI in Azure: Content Safety, Transparency Notes, impact assessments.

■ BUILD (60 MIN)

Take the official Azure AI-102 practice assessment:

learn.microsoft.com/credentials/certifications/azure-ai-engineer/practice/assessment

Target: 80%+ before booking the real exam

Review every wrong answer carefully.

Book your exam for Day 147 at Pearson VUE.

Passing score: 700/1000. Duration: 2 hours.

✓ DONE WHEN: Practice assessment taken (target 80%+). Exam booked for Day 147.

DAY 140 · PORTFOLIO WEBSITE

A professional home for all your work.

■ INDUSTRY (15 MIN)

Browse: ML engineer personal websites. Note what works and what doesn't.

■ STUDY (30 MIN)

A portfolio website needs 5 things: name, what you do, 3 projects, certifications, contact.

That's it. No more. Simple is professional. Complex is noise.

■ BUILD (2.5 HRS) — Portfolio website

Build a single-page portfolio website and deploy it for free.

Options: GitHub Pages (free), Vercel (free), Netlify (free).

Content:

Header: your name + 'AI Engineer' + location

About: 3 sentences. Who you are, what you build, what you're looking for.

Projects: 3 cards with screenshot/GIF, description, and live link

Certifications: list all earned credentials with dates

Contact: GitHub link, LinkedIn link, email

This URL goes on every job application.

✓ DONE WHEN: Portfolio website live at a public URL. All sections complete. URL ready for applications.

WEEK 3 · DAYS 141–150 · INTERVIEW PREP + JOB ATTACK

DAY 141 · RESUME + LINKEDIN

Your application materials.

■ INDUSTRY (15 MIN)

Search: 'ML engineer resume tips what hiring managers look for 2025'

■ STUDY (60 MIN) — Resume best practices

ML engineer resume: 1 page. No exceptions.

Structure: name/contact/links → summary → skills → experience/projects → certifications

Summary: 2 sentences. 'AI engineer with [X] months of project experience in...'

Projects section is more important than work experience if you're career changing.

Each project: one line problem statement + one line on what you built + one metric.

Skills: Python, PyTorch, LangChain, RAG, fine-tuning, Docker, AWS, Azure

■ BUILD (2 HRS) — Resume + LinkedIn

Write your 1-page resume. Save as PDF.

Update LinkedIn:

Headline: 'AI Engineer | LLM Fine-tuning | Agentic AI | [Top Cert]'

About section: expand your resume summary to 5 sentences

Projects section: add your 3 portfolio projects with links

Certifications: add every credential earned in 150 days

Open to work: turn on if you're ready to be found

✓ DONE WHEN: 1-page resume complete. LinkedIn updated with all projects and certifications.

DAY 142 · TECHNICAL INTERVIEW PREP — ML FUNDAMENTALS

Know your foundations cold.

■ INDUSTRY (15 MIN)

Search: 'ML engineer interview questions examples 2025'

■ STUDY (2 HRS) — Core ML interview topics

Practice explaining these out loud (not just in your head):

Bias-variance tradeoff: what it is, how to diagnose, how to fix

Gradient descent variants: SGD vs Adam vs AdamW — when to use each

Overfitting: what causes it, 5 techniques to prevent it

Precision vs recall: the tradeoff, when each matters more

Cross-validation: why it's better than a single train/test split

ROC-AUC: what it measures, when it misleads

Attention mechanism: explain it in 3 sentences to a non-ML engineer

■ BUILD (45 MIN) — ML concepts document

Write your own explanation of each concept above.

No copying from Wikipedia. Write what you actually understand.

If you can't write it, you can't explain it in an interview.

Commit as ML_CONCEPTS.md in your repo.

✓ DONE WHEN: All 7 ML concepts explained in writing. ML_CONCEPTS.md committed.

DAY 143 · SYSTEM DESIGN INTERVIEW PREP

Design ML systems at scale.

■ INDUSTRY (15 MIN)

Read: Chip Huyen's ML systems design introduction — free online.

This is the reference every ML system design interviewer is thinking from.

→ huyenchip.com/machine-learning-systems-design/toc.html

■ STUDY (2 HRS) — Common system design questions

Practice designing these systems (5 minutes plan, 15 minutes explain):

Build a real-time fraud detection system for an e-commerce platform

Build a recommendation system for a video streaming service

Build a content moderation system for a social network

Build an LLM-powered customer support system at scale

For each: data pipeline, model choice, serving architecture, monitoring.

■ BUILD (45 MIN) — System design documents

Write a 300-word system design for ONE of the above problems.

Include: data sources, model type, training pipeline, serving infrastructure,

monitoring approach, failure modes, and how you'd scale to 10x traffic.

Commit as SYSTEM_DESIGN.md

✓ DONE WHEN: System design document written for one problem. All components addressed. 300 words.

DAY 144 · CODING INTERVIEW PREP

Algorithms and data structures for ML roles.

■ INDUSTRY (15 MIN)

Search: 'what coding problems do ML engineers get asked in interviews'

■ STUDY (90 MIN) — Common ML coding questions

ML coding interviews are different from software engineering interviews.

Common question types:

Implement a neural network forward/backward pass from scratch

Implement gradient descent for linear regression

Implement k-means clustering

Implement cross-validation from scratch

Implement precision, recall, F1 score calculation

Debug a training loop that isn't converging

■ BUILD (90 MIN) — Implement 3 from memory

Without looking at any code: implement 3 of the above in 90 minutes.

Gradient descent for linear regression

K-means clustering

Precision, recall, F1 calculation

Test each implementation. Fix until it works.

If you can implement these from memory, you can do it in an interview.

✓ DONE WHEN: 3 ML algorithms implemented from memory and tested without looking at reference code.

DAY 145 · BEHAVIORAL INTERVIEW PREP

Tell your story clearly.

■ INDUSTRY (15 MIN)

Search: 'how to answer tell me about yourself for AI engineer roles'

■ STUDY (60 MIN) — STAR method for ML

STAR: Situation, Task, Action, Result.

Prepare answers to these questions:

'Tell me about yourself.' (2 minutes, no longer)

'Why do you want to work in AI?'

'Tell me about a project you're most proud of.'

'Tell me about a time you failed and what you learned.'

'Where do you see AI going in the next 5 years?'

Your story: zero to AI engineer in 150 days is a compelling narrative. Use it.

■ BUILD (90 MIN) — Write and practice all 5 answers

Write out your answer to each question.

Then: record yourself saying them. Watch the recording.

Notice: do you say 'um' a lot? Do you look away? Is your answer under 2 minutes?

Practice until your 'Tell me about yourself' is smooth, clear, and under 90 seconds.

✓ DONE WHEN: All 5 behavioral questions written and practiced out loud. Recording reviewed.

DAY 146 · COMPANY RESEARCH + TARGET LIST

Know who you're applying to.

■ INDUSTRY (15 MIN)

Browse: LinkedIn Jobs — search 'AI engineer' or 'ML engineer'. Look at what's out there.

■ STUDY (90 MIN) — Build your target list

Research 15 companies you'd actually want to work for.

For each company, find out:

What AI products do they have or build?

Have they published any AI research or engineering blog posts?

What tech stack do they use (matches your skills?)

What ML/AI roles are currently open?

Narrow to your top 10. These get applications on Days 148–149.

→ **levels.fyi/t/machine-learning-engineer**

■ BUILD (60 MIN) — Application materials per company

For each of your 10 target companies:

Find the specific job listing URL

Write one sentence about why you want to work there specifically

Identify one thing about your portfolio most relevant to their work

Customize your resume summary line for this company

Personalization takes 10 minutes per application and doubles response rates.

✓ DONE WHEN: 15 companies researched. Top 10 selected. Customized materials ready for each.

DAY 147 · AZURE AI-102 EXAM DAY

Your fifth and final certification.

■ MORNING PREP (60 MIN)

Read through your Azure AI-102 notes one final time.

Focus on the NLP domain — it's 30–35% of the exam.

Key services to know cold: Language, Custom Vision, Document Intelligence,

Azure OpenAI, Bot Framework, Content Safety.

■ SIT THE AZURE AI-102 EXAM

Online proctored from home via Pearson VUE.

2 hours. Passing score: 700/1000.

If you pass: screenshot immediately. Update GitHub README and LinkedIn.

If you need to retake: book again within 24 hours. Study Domain 4 (NLP).

→ learn.microsoft.com/en-us/credentials/certifications/azure-ai-engineer/

■ AFTER THE EXAM (90 MIN)

Update your GitHub profile README with Azure AI-102 certification.

Update your LinkedIn with the new certification.

Update your resume PDF.

Update your portfolio website.

You now have 5 industry certifications. List them all prominently.

✓ DONE WHEN: Azure AI-102 exam attempted. Credential updated across GitHub, LinkedIn, resume, and portfolio.

DAY 148 · SEND APPLICATIONS 1–5

Day 148. Applications go out.

■ INDUSTRY (15 MIN)

Search: 'how to follow up after submitting a job application AI jobs'

■ PREP (30 MIN) — Final materials check

Verify: portfolio website loads. All project links work. Resume PDF is correct.

Verify: GitHub profile is updated. LinkedIn is updated.

Read your application for Company 1 one final time.

■ APPLY (2.5 HRS) — Applications 1–5

Apply to your top 5 companies. For each application:

Attach your resume PDF

Include your portfolio website URL in every application

Use your customized summary line for this company

Write a 3-sentence cover letter: why AI, why this company, what you built

Screenshot confirmation of submission

After each application: commit the screenshot to a folder called 'applications/' in GitHub.

✓ DONE WHEN: Applications 1–5 submitted. Confirmation screenshots committed to GitHub.

DAY 149 · SEND APPLICATIONS 6–10

Five more. Then you're done.

■ INDUSTRY (15 MIN)

Search: 'AI job market outlook 2025 2026' — where the opportunities are growing.

■ APPLY (3 HRS) — Applications 6–10

Apply to companies 6–10. Same process as Day 148.

Don't rush. Each application gets your full attention.

After submitting all 10: take a screenshot of all 10 confirmation emails.

Commit to GitHub as final proof of completion.

✓ DONE WHEN: Applications 6–10 submitted. All 10 confirmation screenshots committed to GitHub.

DAY 150 · DAY 150 — YOU ARE AN AI MASTER

The finish line. And the starting line.

■ REFLECT (60 MIN)

Open your Day 0 log entry. Read it.

Open your nullvector-journey GitHub repo. Count the green squares.

Read your first commit. Read your last commit.

That distance is real. You created it.

Write your Day 150 log entry. It should be your longest one.

Answer: What were you most afraid of on Day 0?

What surprised you most about the journey?

What would you tell the person who was on Day 0?

What are you building next?

■ 150-DAY RECORD

What you earned in 150 days:

Kaggle Learn: 5 certificates

Harvard CS50 AI: 14 projects + certificate

NVIDIA NCA-AIIO: AI Infrastructure certification

AWS ML Specialty MLS-C01: Machine Learning certification

Azure AI-102: AI Engineer Associate certification

GitHub: 150 consecutive green squares

Portfolio: 10 live AI applications at public URLs

Demo video: your milestone agent in action

10 job applications sent with full portfolio

What you built:

Phase 1: Python foundation, 5 Kaggle certs

Phase 2: CS50 AI — 14 Harvard projects

Phase 3: ML classifier, sentiment analyzer, custom image classifier

Phase 4: Custom image CNN, style transfer, deep learning capstone

Phase 5: Mini-GPT, language model deployed

Phase 6: Fine-tuned LLM, RAG production app

Phase 7: Autonomous agent, multi-agent system, MLOps pipeline

✓ DONE WHEN: Day 150 complete. 10 applications sent. 5 certifications earned. 150 commits. AI master.

DAY 150 MILESTONE · AZURE AI ENGINEER ASSOCIATE + 10 JOB APPLICATIONS SENT

Day 150. You Did It.

Five industry certifications. Ten live AI applications. One hundred and fifty consecutive days of GitHub commits. A portfolio that tells a clear story: I started from zero. I showed up every day. Here is what I built. The applications are sent. The work is visible. The skills are real. What happens next is up to you. But you are no longer someone who wonders if they could do this. You are someone who did.