



AI IN 7 DAYS

DAY 7 HANDBOOK - SELF-STUDY GUIDE

Production, Hardware, Security and Responsible AI

6 VIDEOS	60 MINUTES	18 CANVAS SECTIONS	1 FINAL PITCH
--------------------	----------------------	------------------------------	-------------------------



© 2026 AI PANTHIYA

අද ඔබ ඉගෙනගන්නේ

- Experiment එකක් reliable production AI system එකක් බවට ගෙනයන lifecycle එක
- Deployment, containers, environments, CI/CD, release සහ rollback concepts
- MLOps, LLMOps, testing, registries, versioning සහ change control
- Logs, metrics, traces, monitoring, scaling, drift සහ incident response
- CPU, GPU, cloud, edge, TinyML සහ efficient inference trade-offs
- Security threats, privacy, fairness, governance සහ human oversight
- Final AI Solution Canvas + Risk Register + 60-second presentation



Day 7 Main Task: Final AI Solution Canvas + Risk Register + 60-Second Presentation එක complete කිරීම.

HOW TO USE THIS HANDBOOK

Day 7 එක final integration day එකයි. Day 1 සිට Day 6 දක්වා ඉගෙනගත් problem, data, models, application design සහ evaluation concepts එකට connect කරලා, practical real-world AI solution එකක් responsibly deploy සහ manage කරන framework එක අද build කරනවා.



Day 7 Learning Outcomes

- Notebook result එක සහ production service එක අතර difference පැහැදිලි කිරීම
- Development, Staging සහ Production environments roles differentiate කිරීම
- Containers, Git branches, Pull Requests, CI/CD, release සහ rollback basic flow එක තේරුම්ගැනීම
- MLOps/LLMOps context එකේ data, model, prompt, index සහ config versioning වටහාගැනීම
- Monitoring plan එකක්: logs, metrics, traces, latency, cost, drift, incidents
- Hardware හා deployment options compare කිරීම: CPU/GPU, cloud/local/edge
- Security, privacy, governance සහ meaningful human oversight controls design කිරීම
- 18-section Final AI Solution Canvas complete කිරීම

Today's 6 Video Map

1. From Experiment to Production System
2. MLOps, LLMOps, Testing and Versioning
3. Monitoring, Observability, Scaling and Drift
4. AI Hardware, Cloud, Edge and Efficient AI
5. AI Security, Privacy and Responsible Governance
6. Final AI Solution Canvas and Learning Roadmap

Core rule: Good model score එකක් තිබුණා කියලා alone production-ready solution එකක් වෙන්නේ නැහැ. **software + infrastructure + testing + monitoring + security + governance + human responsibility** එකට වැඩි කළාමයි system එක truly usable වෙන්නේ.

From Experiment to Production System

Notebook එකක result එක real usersට safe, stable service එකක් වෙන්න තවත් layers ගොඩක් අවශ්‍යයි.

Notebook → Production Lifecycle

Experiment



Test



Package



Deploy



Monitor



Maintain

Production System Needs

- Interface – user interaction layer
- Dependencies – libraries, packages, runtime
- Testing – expected behaviour verify කිරීම
- Security – data, users, access protection
- Deployment – real environment publish කිරීම
- Monitoring – errors, latency, uptime track කිරීම
- Support – incidents handle කරන process
- Versioning + Rollback – safe change management

Key Concepts

- **Training-Serving Mismatch:** training preprocessing සහ production preprocessing වෙනස් වීම
- **Best practice:** preprocessing + model එක complete pipeline එකක් ලෙස save කරන්න
- **Model Artifact:** model parameters, preprocessing logic, config, metadata, labels, dependencies
- **Container:** app + runtime + dependencies bundle කිරීම
- **Environment split:** Development → Staging → Production

Development

Build, debug, experiment කරන place. Fast iteration සඳහා.

Staging

Production-like testing area. Release එක real usersට පෙර validate කරනවා.

Production

Real users සේවය ලබන live environment එක. Reliability, security සහ rollback readiness critical.

Remember: Container use කළා කියලා automatically security, scalability හෝ correctness guarantee වෙන්නේ නැහැ.

MLOps, LLMOps, Testing and Versioning

MLOps සහ LLMOps කියන්නේ single tool එකක් නෙවෙයි. ඒවා repeatable, auditable operations practices එකතුවක්.

Basic Release Flow



What to Version?

- Code and configuration
- Training / reference data
- Model checkpoints and artifacts
- Prompts and prompt templates
- Embeddings / vector indexes
- Evaluation datasets and test cases
- External tool settings / connectors

CI vs CD

CI: changes integrate කරලා tests run කිරීම. **CD:** approved build safely release කිරීම.

Registry

Models, prompts, datasets හෝ artifacts official version catalog එකකින් manage කිරීම.

Unit Test

Small functions or components isolated ලෙස work කරනවද බලන test.

Integration Test

Multiple components එකට connect වුණම flow එක හරියට work කරනවද බලන test.

Regression Test

New release එක critical known behaviour break කරලද කියලා check කිරීම.

Rollback

Issue එකක් ආවම known safe previous version එකට ආපසු යාම.

Why version everything? Reproduce කිරීම, compare කිරීම, audit trails තබාගැනීම, exact system state එකට rollback කිරීම සහ change impact තේරුම්ගැනීම සඳහා.

Monitoring, Observability, Scaling and Drift

Production system එක launch කරලා ඉවර නෙවෙයි. Operation quality දිගටම measure කරලා incidents detect කරන්න ඕන.

Logs

Individual events record කරනවා. Example: request received, auth failed, tool timeout.

Metrics

Aggregated numbers. Example: latency, success rate, cost per request, uptime.

Traces

End-to-end request journey. Example: user query → retrieval → model → tool call → final answer.

What to Monitor?

- Latency, throughput, concurrency
- Error rate, timeout rate, retry rate
- Quality metrics and user feedback
- Cost per task and budget consumption
- Tool-call failures and prompt traces
- Security alerts and permission denials

Drift Types

- **Input Drift:** incoming data distribution වෙනස් වීම
- **Data Drift:** reference/training data compared to new data shift වීම
- **Concept Drift:** input-output relationship itself වෙනස් වීම
- **Model Decay:** visible performance drop over time

Incident Response Loop

Alert



Triage



Contain



Fix / Rollback



Communicate



Learn

Important: Monitoring හැම failure එකක්ම හෝ drift එකේ exact cause එකම automatically හඳුනාගන්නේ නැහැ. Human investigation තවමත් අවශ්‍යයි.

AI Hardware, Cloud, Edge and Efficient AI

Every workload එකකට එකම hardware choice එක හොඳම නෙවේ. Task, latency, privacy, cost සහ scale එක අනුව decide කරන්න.

CPU

General-purpose. Small workloads, orchestration, classical ML, cost-sensitive cases සඳහා often enough.

GPU

Parallel computation strong. Training, embeddings, deep learning inference සඳහා useful.

TPU / NPU

Specific accelerator types. Some platforms/devices වල efficient AI workload support කරයි.

Cloud

Elastic scale, managed services, quick access. But ongoing cost සහ data-location/privacy concerns consider කරන්න.

Local / On-Prem

More control and possible privacy advantages. Setup, maintenance, capacity responsibility ඔබ සතුව.

Edge / TinyML

Low latency, offline use, device-side inference. But model size and compute constraints stronger.

Deployment Trade-off Matrix

- **Cloud:** flexible scale, faster setup
- **Local:** more direct control, possible data sensitivity fit
- **Edge:** low latency, offline operation
- **Question:** Which option matches privacy, cost, connectivity and scale needs?

Efficient Inference Techniques

- **Quantisation:** lower numerical precision use කිරීම
- **Pruning:** less important parameters remove කිරීම
- **Distillation:** larger model behaviour smaller model එකකට transfer කිරීම
- **Batching / Caching:** repeated or grouped work efficient කිරීම

Note: Quantisation, pruning සහ distillation speed/cost reduce කළත් quality එකට impact වෙන්න පුළුවන්. Always test before production use.

AI Security, Privacy and Responsible Governance

Useful AI system එකක් safe, fair සහ accountable system එකක් විය යුතුයි.

Prompt Injection

User හෝ retrieved content malicious instructions දී system behaviour manipulate කිරීම.

Data Leakage

Sensitive or unauthorised information accidentally expose වීම.

Poisoning

Training or index data corrupt කර model/retrieval behaviour harm කිරීම.

Adversarial Examples

Input slightly change කර model confuse කිරීම.

Model Theft

Model behaviour, weights හෝ valuable outputs extract කිරීම.

Supply-Chain Risk

Third-party packages, models, APIs හෝ tools හරහා risk introduce වීම.

Core Controls

- Least privilege access
- Secrets management
- Input/output validation
- Approval gates for high-risk actions
- Audit logs and traceability
- Data minimisation and privacy protection
- Fallbacks and safe failure modes

Governance Lifecycle

Govern - roles, policies, responsibilities



Map - use case, data, context, stakeholders



Measure - risk, quality, fairness, robustness



Manage - controls, incidents, updates, retirement

Meaningful Human Oversight

Human reviewerට correct decision ගන්න **information + competence + time + authority + intervention power** තිබිය යුතුයි. Decorative checkbox එකක් පමණක් නම් එය real oversight නෙවේ.

When NOT to Use AI

Simple rule-based solution enough නම්, good data නැත්නම්, harm unacceptable නම්, output verify කළ නොහැකි නම්, cost value එකට වඩා වැඩි නම්, හෝ accountability loss high නම් AI use නොකළ හැක.

Final AI Solution Canvas

Days 1-7 සියල්ල එකතු කර practical, responsible AI solution plan එකක් පුරවන්න.

18-Section Final AI Solution Canvas

1. Problem

What real pain are you solving?

2. Intended Users

Who uses or is affected?

3. Value & Success

Measurable benefit and success evidence?

4. Why AI?

Why not a simpler deterministic solution?

5. Input Data

What data/documents/signals are needed?

6. Data Quality & Permission

Source, quality, coverage, consent, ownership?

7. Task & Method

Prediction, retrieval, generation, rules, recommendation or hybrid?

8. Model or Service

Which algorithm, API or model and why?

9. Application Architecture

Frontend, backend, storage, API, authentication?

10. Retrieval & Tools

Search, RAG, memory, tools, permissions?

11. Hardware & Deployment

CPU/GPU, cloud/local/edge, scale, cost?

12. Evaluation

Quality, user value, robustness, latency, cost?

13. Security & Privacy

Threats, secrets, access controls, protection?

14. Human Approval & Accountability

Who reviews, overrides and remains responsible?

15. Monitoring

Quality, drift, incidents, cost and owners?

16. Limitations & Risks

Bias, hallucination, misuse, wider impacts?

17. Rollback / Fallback

What happens on outage, unsafe input or failure?

18. Responsible Owner

Who maintains, approves updates and retires the system?

Canvas rule: Completed canvas එක useful planning document එකක්. ඒත් safety හෝ legal compliance automatic ලෙස prove කරන්නේ නැහැ. Validation, testing සහ review තවමත් අවශ්‍යයි.

RISK REGISTER • 60-SECOND PITCH • LEARNING ROADMAP

Threat / Control Register

Risk / Threat	Control	Owner
Prompt injection via external content	Tool scope limits, source filtering, answer grounding	Product + Engineering
Sensitive data leakage	Access control, redaction, logging, review	Security Lead
Bad release quality	Regression tests, staging, approval gate, rollback	ML / Ops
Cost spike	Budget alerts, caching, model routing	Ops / Finance
Incorrect high-impact output	Human approval + fallback path	Responsible Owner

60-Second Presentation Template

1. **Problem & Users:** What problem? For whom?
2. **AI Solution:** workflow, data and model in one simple flow.
3. **Value & Evaluation:** what evidence shows it works?
4. **Main Risk:** one major limitation or threat.
5. **Human Control:** who reviews and remains accountable?
6. **Next Step:** prototype, test or validation action.

Path 1: Builder

RAG, vector DBs, agents, tool calling, multimodal apps සහ evaluation deeper build skills develop කරන්න.

Path 2: Operations

CI/CD, versioning, registries, monitoring, incidents, governance සහ production operations focus කරන්න.

Path 3: Responsible AI

Safety, privacy, fairness, explainability, governance, regulation සහ human oversight stronger කරන්න.

30-Day Goal Rule: “Everything එකවර ඉගෙනගන්න” වෙනුවට measurable next step එකක් තෝරන්න. Example: **Next 30 days වල complete RAG Study Assistant prototype එකක් build කර documentation එකක් publish කරනවා.**

DAY 7 QUICK SELF-CHECK

Quiz Questions

1. Notebook සහ production system අතර difference එක මොකක්ද?
2. Container guarantee නොකරන දේවල් දෙකක් කියන්න.
3. Development, staging සහ production roles explain කරන්න.
4. CI සහ CD අතර difference එක?
5. Rollback කියන්නේ මොකක්ද?
6. Logs, metrics සහ traces වෙනස් වන්නේ කොහොමද?
7. Input drift, concept drift සහ model decay අතර වෙනස?
8. Cloud vs edge deployment choose කරන reasons?
9. Quantisation, pruning සහ distillation concepts?
10. Prompt injection සහ data leakage මොනවාද?
11. Least privilege define කරන්න.
12. Govern-Map-Measure-Manage lifecycle explain කරන්න.
13. AI use නොකළ යුතු reasons තුනක්?
14. Final AI Solution Canvas එකේ section කිහිපයක් list කරන්න.
15. Responsible ownerගේ role මොකක්ද?

Answer Hints

- Experiment result vs tested, secure, operated service.
- Security and scalability automatic නෙවේ.
- Build / validate / live user service.
- Integrate & test changes vs release approved builds.
- Return to known safe version.
- Events / numeric trends / end-to-end request path.
- Data shift / relationship shift / visible performance decay.
- Elastic scale vs privacy, low latency, offline use.
- Lower precision / remove parameters / transfer to smaller model.
- Manipulation instructions / unauthorised data exposure.
- Minimum required permissions only.
- Set governance → map context → measure risk → manage controls.
- Simple rule enough, bad data, unacceptable harm, unverifiable output, high cost.
- Problem, users, data, model, evaluation, monitoring, risks, owner ...
- Maintain, approve changes, respond to incidents and retire system.



Sampath Bandaranayake
Instructor - AI Panthiya

Need guided learning?

Online Zoom Classes - Contact Us on WhatsApp

WhatsApp: +81 70 2313 1981

MSc - AI and Digital Innovation (Final Year) - Tokyo International University, Japan

MBA - Cardiff Metropolitan University, UK

BSc - Applied Sciences, Sri Lanka

LinkedIn: <https://www.linkedin.com/in/sampathbandara/>

Website: <https://aipanthiya.com/>

Course Completion Note: මේ 7-day course එකෙන් ඔබට AI ගැන **AI-literate informed beginner / builder foundation** එකක් ලැබෙනවා. Professional expert qualification එකක් නෙවෙයි. දැන් practical projects, evaluation සහ responsible improvement තුළින් deeper යන්න.