

The AI Production Readiness Checklist

Pre-Production

Before Go-Live

- ☒ Is there a prompt governance standard documented and shared across the team?
- ☒ Is AI output being actively reviewed or blindly accepted?
- ☒ Is there a knowledge base capturing what AI has built and why?
- ☒ Is traceability documented for every AI-generated module?
- ☒ Is there a tested fallback when the AI layer fails?
- ☒ Has the team been trained on the specific model's known failure modes?
- ☒ Is there an observability layer monitoring AI behaviour in production?
- ☒ Has a human-in-the-loop review been defined for high-stakes outputs?

Post-Production

Month 3 and Month 6

- ☒ Can any engineer explain what AI built and why, without AI assistance?
- ☒ Is prompt intelligence captured in a shared system, not individual heads?
- ☒ Are DORA metrics being tracked before and after AI adoption?
- ☒ Is there a governance review scheduled at month 3 and month 6?
- ☒ Does the CXO have real-time visibility into how AI is being used?
- ☒ Is the productivity gain reaching customers, or only individuals?
- ☒ Is the AI layer in your architecture still the right pattern for current scale?
- ☒ Is there a plan for when the AI model you depend on is deprecated?

Engineering visibility and AI governance in one platform.

getpms.cloudbyte.ai