

AI Readiness

Checklist

Score your organization across 8 dimensions to understand where you stand before investing in AI.

Most AI projects fail not because of technology, but because of organizational readiness. This checklist helps you identify gaps before they become expensive problems. Score each dimension honestly — the goal is clarity, not a high score.

01 Data Readiness

Do you have the data to train and feed AI systems?

- Core business data is digitized and accessible via APIs or databases
- Data quality is monitored — you know your error rate, missing fields, and staleness
- You have at least 6 months of historical data for your target use case
- Data ownership and access permissions are clearly defined

02 Infrastructure

Can your systems support AI workloads?

- You have a cloud environment (AWS/Azure/GCP) with GPU or inference capacity
- CI/CD pipelines exist for deploying and updating services
- You can provision new services without a 3-month procurement cycle
- Logging and monitoring infrastructure is in place (Grafana, Datadog, etc.)

03 Team Skills

Does your team have the capability to build and maintain AI systems?

- At least one engineer understands LLM APIs, prompt engineering, or ML pipelines
- Your team can build and deploy Python or TypeScript services independently
- Someone on the team has shipped a data pipeline or ML model to production
- You have a plan for upskilling or hiring AI-capable engineers

Have you identified a specific, high-value problem for AI to solve?

- You can name the specific workflow or process AI will improve
- You can quantify the current cost of the problem (hours, errors, revenue lost)
- The use case has a clear success metric (accuracy, speed, cost reduction)
- Stakeholders agree on the priority and expected outcome

05 PoC-to-Production Path

Can you ship AI from demo to production reliably?

- You have a process for moving from prototype to production deployment
- Testing strategy includes AI-specific concerns (hallucination, drift, edge cases)
- You've defined acceptable error rates and fallback behavior
- There's a timeline and budget for the PoC-to-production transition

06 Monitoring & Observability

Can you detect when AI systems misbehave?

- You plan to monitor AI output quality, not just uptime
- Alerting is configured for anomalies in AI behavior (confidence drops, latency spikes)
- You have a human review process for edge cases and low-confidence outputs
- Feedback loops exist to improve model performance over time

07 Governance & Compliance

Are you prepared for the regulatory and ethical dimensions of AI?

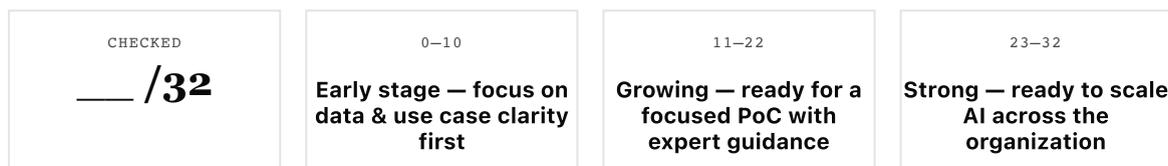
- You understand which regulations apply (GDPR, EU AI Act, industry-specific)
- Data handling policies cover AI training data and model outputs
- There's a process for auditing AI decisions when required
- You've assessed bias risk for your specific use case and data

08 ROI Measurement

Can you prove AI is delivering value?

- You have baseline metrics for the process AI will improve
- Success criteria are defined before implementation begins
- You can measure cost savings, time savings, or revenue impact specifically
- There's a review cadence to evaluate AI ROI (monthly or quarterly)

YOUR SCORE



Want to discuss your score?

I help organizations close the gaps this checklist reveals — from data strategy to production AI systems. Let's have an honest conversation about where you stand.

dvorson.github.io/#contact