



Azure Document Intelligence: Enterprise Deployment & Readiness Checklist

Use this 4-phase technical checklist to evaluate your organization's readiness before initializing Azure AI Classification models. Score your readiness to prevent downstream extraction failures and API throttling.

Phase 1: Architecture & Governance Readiness

Establish the secure foundation and identify the exact business boundaries for your AI pipeline.

- **Workflow Mapping:** We have documented all incoming document types, their current manual triage paths, and the downstream systems they belong to.
- **Business Routing Rules:** We have defined classification boundaries based on operational actions (e.g., "Approval Path A vs. Approval Path B") rather than minor visual differences.
- **Compliance Verification:** Azure Virtual Network (VNet) integration and Role-Based Access Control (RBAC) protocols are approved by the security team for handling sensitive PII/financial data.
- **Destination Architecture:** Target repositories (SharePoint libraries, vector databases for RAG, or Power Automate flows) are configured to receive the extracted metadata.

Phase 2: Data Collection & Sample Engineering

Prevent model brittleness by engineering a highly resilient training dataset.

- **Sample Diversity:** We have collected a minimum of 20-50 samples per document class, representing various vendors, layouts, and historical periods.



- [] **Edge Case Inclusion:** The training set actively includes "ugly" real-world files: low-resolution scans, skewed pages, mobile phone photos, and documents with handwritten notes or stamps.
- [] **Mixed Packet Identification:** We have identified common multi-document packet scenarios (e.g., an invoice combined with a delivery note in a single PDF) to utilize Azure's auto-split functionality.
- [] **Labeling Protocol:** A dedicated subject matter expert (SME), not just a developer, has reviewed and labeled the training data to ensure business accuracy.

Phase 3: Model Configuration & Technical Integration

Connect the classification logic to extraction models and enterprise APIs.

- [] **Composed Model Setup:** We have architected explicit classification within our Azure v4.0 composed models to act as the primary conditional router.
- [] **Split Mode Configuration:** PDF ingestion pipelines are configured with the appropriate split mode (**None**, **perPage**, or **Auto**) based on our mixed-packet audit.
- [] **API Authentication:** REST API or SDK calls are secured using Managed Identities, eliminating hardcoded access tokens in our repository.
- [] **Exception Branching:** We have built a dedicated manual-review loop (a "human-in-the-loop" queue) for documents that fall below the established confidence threshold.

Phase 4: Post-Launch & Cloud Economics

Monitor the metrics that actually impact your bottom line and RAG capabilities.

- [] **Throttling & API Limits:** We have evaluated our peak upload volumes and aligned them with our Azure pricing tier to prevent unexpected API throttling.
- [] **Waste Reduction:** We have a pre-processing step to trim blank or irrelevant pages (e.g., terms and conditions) to reduce cost-per-page billing.



- **Drift Monitoring:** A protocol is in place to retrain the classifier when exception rates climb due to new vendor templates or changes in scanner hardware.
- **LLM Anchoring (Optional/Advanced):** Extracted metadata is cleanly structured and injected with semantically dense phrases to fuel downstream Retrieval-Augmented Generation (RAG) vector embeddings.

Sign-off:

Project Lead: _____

Date: _____

Deployment Status: Ready Needs Revision