



Microsoft Azure AI Engineer Associate Certification Study Notes

Code: AI-102

Azure AI Services



Azure AI Services Overview

Platform for AI solutions

Service Categories

Category	Services	Capabilities
Vision	AI Vision, Custom Vision, Face	Image analysis, object detection, OCR
Language	Language, Translator	NLP, sentiment, translation, QA
Speech	Speech Service	STT, TTS, translation, speaker recognition
Decision	Content Safety, Personalizer	Content moderation, recommendations
OpenAI	Azure OpenAI Service	GPT, DALL-E, embeddings
Document	Document Intelligence	Form extraction, document analysis

Provisioning Options
Analytics (🔗 ⚙️ A)

Feedback

Multi-service

Single resource for multiple AI services, one key

Single-service

Dedicated resource for one service, separate billing

Commitment Tier

Discounted pricing for committed usage

Security & Access

- **Keys:** Primary and secondary keys for rotation
- **Entra ID:** Managed identity, service principals
- **Virtual Network:** Private endpoints, service endpoints
- **RBAC:** Fine-grained access control

Exam Focus Areas

- Multi-service resource simplifies management with one endpoint
- Use managed identity over keys when possible
- Private endpoints keep traffic off public internet
- Key rotation: update apps to use key2, regenerate key1

Computer Vision

Azure AI Vision

Image and video analysis

Image Analysis 4.0

Feature	Description	API
Captions	Generate natural language descriptions	analyze (features=caption)
Dense Captions	Describe multiple regions	analyze (features=denseCaptions)
Tags	Identify objects, concepts	analyze (features=tags)
Objects	Detect with bounding boxes	analyze (features=objects)
People	Detect people in images	analyze (features=people)
Read (OCR)	Extract text from images	analyze (features=read)

Custom Vision

Classification

Categorize whole images into classes

Object Detection

Locate and classify objects

Training

Upload tagged images, train model

Export

ONNX, TensorFlow, CoreML for edge

Face Service

- **Detection:** Find faces, return face rectangles
- **Attributes:** Age, emotion, glasses, blur (requires approval)
- **Verification:** Compare two faces (1:1)
- **Identification:** Match face to person group (1:N)
- **Limited Access:** Identification requires approval

Limited Access Features

Face identification, celebrity recognition, and some attributes require Microsoft approval. Apply via Limited Access form.



Exam Focus Areas

- Image Analysis 4.0 uses Florence foundation model
- Custom Vision for domain-specific image classification
- Face API: Detection is open, Identification needs approval
- Use SDK or REST API to call vision services

Language Services



Azure AI Language

Natural language processing

Pre-built Features

Feature	Description	Output
Sentiment Analysis	Positive/negative/neutral + scores	Document and sentence level
Key Phrase Extraction	Main topics in text	List of phrases

Feature	Description	Output
Named Entity Recognition	People, places, orgs, dates	Entities with categories
Entity Linking	Link entities to Wikipedia	Entity with URL
Language Detection	Identify text language	Language code + confidence
Summarization	Extractive and abstractive	Summary text

Custom Features

CLU

Conversational Language Understanding - intents and entities

Custom NER

Train to extract domain-specific entities

Custom Classification

Classify text into custom categories

Custom QA

Build from documents, FAQs

Question Answering

- Create knowledge base from URLs, files, or manual QA pairs
- Supports multi-turn conversations
- Active learning suggests new QA pairs
- Chitchat: Add personality with pre-built responses
- Deploy to web app or integrate with Bot Framework



Exam Focus Areas

- CLU replaces LUIS for intent/entity recognition
- Sentiment returns per-document and per-sentence scores
- NER categories: Person, Location, Organization, DateTime, etc.
- Custom models need labeled training data

Speech Services



Azure Speech Service

Speech recognition and synthesis

Speech Capabilities

Feature	Description	Use Case
Speech-to-Text	Transcribe audio to text	Transcription, voice commands
Text-to-Speech	Generate speech from text	Voice assistants, accessibility
Speech Translation	Translate spoken language	Real-time translation
Speaker Recognition	Identify speakers by voice	Authentication
Pronunciation Assessment	Evaluate pronunciation	Language learning

Speech-to-Text Options

Real-time

Streaming transcription from microphone

Batch

Transcribe audio files asynchronously

Custom Speech

Train on domain vocabulary, accents

Conversation

Multi-speaker diarization

Text-to-Speech Options

- **Neural Voices:** Natural-sounding synthesis
- **Custom Neural Voice:** Create unique voice (requires approval)
- **SSML:** Control pronunciation, pauses, emphasis
- **Audio Profiles:** Optimize for headphones, speakers, phone

SSML Example

Speech Synthesis Markup Language controls speech output:

```
<speak><voice name="en-US-JennyNeural">Hello <break time="500ms"/>
world!</voice></speak>
```



Exam Focus Areas

- Speech SDK supports real-time and batch scenarios
- Custom Speech improves accuracy for domain terminology
- SSML provides fine-grained control over speech output
- Speaker verification: text-dependent or text-independent

 Feedback

Document Intelligence



Azure AI Document Intelligence

Extract information from documents

Pre-built Models

Model	Document Type	Extracted Data
Read	Any document	Text, structure, language
Layout	Any document	Tables, selection marks, structure
Invoice	Invoices	Vendor, dates, amounts, line items
Receipt	Receipts	Merchant, date, total, items
ID Document	IDs, passports	Name, DOB, expiration, address
Business Card	Business cards	Name, company, contact info

Custom Models

Custom Template

Fixed layout documents (forms)

Custom Neural

Variable layouts, semi-structured

Composed Model

Combine multiple custom models

Document Studio

Label and train custom models

Processing Flow

- **1. Submit:** Send document URL or base64 to analyze endpoint
- **2. Poll:** Check operation status until complete
- **3. Results:** Get extracted fields, tables, key-value pairs
- Supports PDF, JPEG, PNG, TIFF, BMP, HEIF



Exam Focus Areas

- Pre-built models work out of the box, no training
- Custom template: minimum 5 labeled samples
- Custom neural: handles variations better than template
- Composed models route to correct sub-model automatically

Azure OpenAI

<> Azure OpenAI Service

Enterprise generative AI

Available Models

Model	Capability	Use Case
GPT-4o	Multimodal reasoning	Complex tasks, vision+text
GPT-4	Advanced reasoning	Complex analysis, coding
GPT-3.5 Turbo	Fast text generation	Chatbots, content
DALL-E 3	Image generation	Create images from text
Whisper	Speech transcription	Audio to text
Embeddings	Text to vectors	Search, similarity

Deployment Options

- Standard**

Pay-per-token, shared capacity
- Provisioned**

Reserved throughput (PTU)
- Global**

Automatic region routing
- Data Zone**

Keep data in geography

Key Concepts

- **Tokens:** Input + output tokens count toward limits
- **System Message:** Set assistant behavior and constraints

- **Temperature:** 0 = deterministic, 1+ = creative
- **Content Filters:** Built-in safety for hate, violence, etc.

On Your Data

Azure OpenAI On Your Data enables RAG without custom code. Connect to Azure AI Search, upload files, or link blob storage. Responses grounded in your documents.



Exam Focus Areas

- Azure OpenAI = OpenAI models + Azure security/compliance
- Content filtering enabled by default, can customize
- On Your Data provides built-in RAG capability
- Use embeddings + vector search for semantic retrieval

Responsible AI



Responsible AI Practices

Building ethical AI solutions

Microsoft Responsible AI Principles

Principle	Description
Fairness	AI should treat all people fairly
Reliability & Safety	Perform reliably and safely
Privacy & Security	Secure and respect privacy

Principle	Description
Inclusiveness	Empower everyone
Transparency	Be understandable
Accountability	People accountable for AI

Content Safety

Categories

Hate, violence, sexual, self-harm

Severity Levels

Safe, Low, Medium, High

Actions

Allow, block based on threshold

Custom Lists

Block specific terms

Implementation

- **Content Filters:** Built into Azure OpenAI, configurable
- **Groundedness Detection:** Check if outputs match source
- **Prompt Shields:** Detect jailbreak attempts
- **Azure AI Content Safety:** Standalone moderation service



Exam Focus Areas

- Know all 6 responsible AI principles

 Feedback

- Content Safety API works with text and images
- Azure OpenAI has content filtering enabled by default
- Limited Access features require application approval

IT certification prep: [Practice questions](#) / [Full exams](#) / [Study guides](#)



CertStud

Free IT certification practice exams and study materials.

A [Filantus](#) product



Resources

Practice Tests

Free IT Practice Tests

Cloud Practice Tests

Cybersecurity Practice Tests

Exam Simulator

Cert Quiz

Roadmaps

AI Roadmap

Study Guides

Blog

AI Corner

Newsletter

Company

About

How it works

Contact

FAQ

Plans & pricing

Filantus

Our Products

Foci

Legal

Privacy Policy

Terms of Service

Focus & productivity

CollegeDecider

College planning & SAT tutoring

BoostLogik

SEO & AEO solutions

WanderingHermit

AI travel planning

Brakto

Tournament management

© 2026 CertStud. All rights reserved.



Affiliate Disclosure: We may earn commissions from qualifying purchases through affiliate links.
[Learn more](#)