

Microsoft Foundry is the application platform for Azure AI

Microsoft Foundry provides one environment for building, evaluating, deploying, and operating AI applications and agents.

It brings together

- Models
- Agents
- Tools and knowledge
- Evaluations and observability
- Azure security and governance

Key message

Business application → Microsoft Foundry → Azure platform and infrastructure.

The Foundry resource provides the governance boundary

A Foundry resource is the top-level Azure resource for the shared AI environment.

A resource controls

- Security and networking
- Model deployments
- Connections
- Access control
- Shared governance

Azure hierarchy

- Tenant and subscription: Organization and billing boundary
- Resource group: Lifecycle grouping for Azure resources
- Foundry resource and projects: Environment governance, then solution work

A resource can contain multiple projects.

Exam callout

A Foundry resource governs the environment. A project organizes development work.

Projects isolate application development

A project is the working boundary for a specific AI solution.

Projects organize

- Agents
- Files and project assets
- Evaluations
- Application development activity
- Project-level access

One Foundry resource

- Customer support: The reference support assistant
- Document processing: Extraction and review workflows
- Employee assistant: Internal knowledge and tasks

Remember

Shared controls stay at the Foundry-resource level. Each project isolates one solution.

Models become usable through deployments

Model

- The underlying AI capability selected from the Foundry Models catalog
- Choosing it does not automatically make it available to an application

Deployment

- A callable instance of that model
- An endpoint
- Capacity and quota
- Regional configuration
- Content-safety controls
- Usage-based or provisioned billing

Exam callout

Model selection answers, “Which capability do we need?” Deployment answers, “How will the application access that capability?”

Foundry Models catalog



Deployment endpoint



Application

Tools and agents extend model capabilities

Tools give an application access to external capabilities, data, or actions. Agents combine a model with instructions, tools, and knowledge to complete a task.

Tools

- Search and retrieval
- Functions and APIs
- Code execution
- Connected business systems

Support assistant

- Model: Interprets the customer's request
- Retrieval tool: Finds relevant support information
- Agent: Decides which information or action the request requires

Remember

A model generates or interprets content. A tool retrieves or acts. An agent decides how to use both.

Portal experimentation and application development serve different stages

Foundry portal

Explore available models

Test prompts in playgrounds

Compare model behavior

Configure prompt agents

Run quick evaluations

Application development

Build repeatable application logic

Call deployments through SDKs or REST APIs

Integrate identity, data, and tools

Add testing, monitoring, and deployment workflows

Operate the solution at production scale

Stage, not substitute

Portal experiments inform the application. The playground is not the finished customer solution.

The complete mental model

Trace an AI solution from the Azure foundation to the experience the customer sees.

Remember this sequence

- Azure provides the cloud foundation.
- A Foundry resource provides shared governance.
- A project organizes one AI solution.
- A model deployment provides an accessible endpoint.
- Tools connect the solution to information and actions.
- An agent coordinates the model and tools.
- The application delivers the user experience.

Key message

Knowledge check: experiment with prompts and compare models in a playground in the Microsoft Foundry portal—before writing application code.