

IT Setup Instructions for Microsoft Foundry Participants

These instructions outline what your IT team may need to provision or enable if participants are building custom AI agents and applications using **Microsoft Foundry** and encounter setup, access, deployment, model, or security restrictions.

This guidance is intended for:

- developer-focused AI projects,
- advanced agent orchestration,
- custom AI applications,
- model experimentation,
- enterprise integrations,
- retrieval-augmented generation (RAG),
- and API-based AI solutions.

This is typically more advanced than:

- Agent Builder in Microsoft 365 Copilot
- or low-code Copilot Studio scenarios.

1. What participants need

Participants typically need:

- A work or school Microsoft Entra ID account
- Access to a Microsoft Azure subscription
- Access to Microsoft Foundry
- Permission to create or use AI services
- Access to a Foundry project/workspace
- Access to approved AI models (such as GPT models)

Depending on the scenario, participants may also require:

- AI Search services
 - Storage services
 - API permissions
 - networking approvals
 - vector database/query services
-

2. What IT may need to provision

Microsoft Foundry access

Users need access to:

- Microsoft Foundry
- a Foundry project/workspace
- associated AI resources

IT may need to:

- provision a Foundry workspace/project
- assign project permissions
- add users to approved security groups

Without project access, participants may not be able to build or run solutions.

Azure subscription access

Participants need access to an Azure subscription with permission to use AI services.

IT may need to:

- assign Azure RBAC roles
- add users to subscriptions/resource groups
- confirm AI services are permitted by organisational policy

Without subscription access, users may not be able to provision or consume AI resources.

AI model access

Many Foundry scenarios require access to approved AI model deployments.

IT may need to:

- enable approved model access
- deploy required models
- assign permissions to model deployments
- confirm regional/model availability

Some organisations restrict:

- model families,
- model regions,

- or deployment types.
-

Appropriate permissions and RBAC roles

Participants may require roles such as:

- Contributor
- Cognitive Services User
- Azure AI Developer
- Storage Blob Data Contributor
- Search Service Contributor

Required roles depend on:

- AI models used,
 - data access,
 - orchestration,
 - storage,
 - and deployment scenarios.
-

Networking and security approvals

Some organisations restrict AI services through:

- private networking,
- Conditional Access,
- firewalls,
- VPNs,
- approved regions,
- or security/compliance policies.

If users cannot access services, IT may need to confirm:

- approved network access
 - browser/security compatibility
 - permitted service regions
 - AI services are not blocked by policy
-

3. Common issues and what participants should ask IT

Users cannot access Microsoft Foundry

Ask IT to confirm:

- subscription access is assigned
 - Foundry access is enabled
 - users are added to the correct workspace/project
 - Conditional Access policies are not blocking sign-in
 - the user is signed into the correct Entra ID tenant
-

Users cannot see AI models

Ask IT to confirm:

- approved model access exists
 - models are deployed into the environment
 - users have permission to the deployments
 - quotas/regional restrictions are not blocking usage
-

Users receive “quota exceeded” or throttling errors

Ask IT to confirm:

- sufficient model quota exists
 - sufficient token/request capacity is available
 - workshop participants are not hitting rate limits
 - capacity has been assigned correctly
-

Users cannot upload or ground on internal documents

Ask IT to confirm:

- storage services are enabled
 - AI Search/indexing services are configured
 - upload policies permit document ingestion
 - users have permission to the underlying files
 - governance policies permit organisational data usage
-

Users cannot connect APIs or external systems

Ask IT to confirm:

- outbound networking policies permit access
 - API endpoints are approved
 - service principals/managed identities are configured
 - secrets/key vault permissions are enabled
 - connector governance policies permit usage
-

Users cannot deploy or publish solutions

Ask IT to confirm:

- deployment permissions are assigned
- required Azure resource providers are registered
- hosting/deployment services are provisioned
- organisational governance policies allow deployment

Some organisations restrict:

- external endpoints,
 - public hosting,
 - or external integrations.
-

4. Recommended minimum setup for workshops and hackathons

For the smoothest participant experience, recommended setup includes:

- Pre-created Foundry workspaces/projects
- Pre-assigned participant permissions
- Pre-deployed AI models
- Pre-approved AI quotas/capacity
- Approved networking policies
- Pre-configured storage/search resources (if grounding is required)

This avoids delays caused by:

- RBAC provisioning,
- quota requests,
- model approvals,
- or security reviews.

5. What participants can do once enabled

Once configured, participants can:

- build custom AI agents
- orchestrate multi-step AI workflows
- use advanced AI models
- ground on organisational knowledge
- integrate APIs and external systems
- implement RAG/search patterns
- prototype enterprise AI applications
- integrate with Copilot Studio, Teams, or Microsoft 365 experiences

6. Important distinction – Agent Builder vs Copilot Studio vs Microsoft Foundry

Platform	Best for
Agent Builder	Lightweight Microsoft 365 declarative agents
Copilot Studio	Low-code enterprise agents and workflows
Microsoft Foundry	Developer-focused custom AI applications and advanced orchestration

Participants building with Microsoft Foundry are typically working in a more developer-oriented environment with significantly more flexibility — but also more IT/security dependencies.

Official Microsoft References

- <https://learn.microsoft.com/azure/ai-foundry/>
- <https://learn.microsoft.com/azure/ai-services/openai/>
- <https://learn.microsoft.com/azure/ai-foundry/how-to/create-projects>
- <https://learn.microsoft.com/azure/ai-services/openai/quotas-limits>
- <https://learn.microsoft.com/azure/role-based-access-control/built-in-roles>