

Microsoft Agent 365

Enterprise reference architecture and best practices

A practical guide for architects, identity, security, developer and business leaders to discover, govern, secure, build and measure AI agents across the enterprise.

16

deep-dive pages

4

agent types compared

75

cited sources

PREPARED BY

Carlos Suarez
Solution Engineer
carlossuarez@microsoft.com

Alejandro Lopez
FastTrack Engineer
Alejandro.Lopez@microsoft.com

CONTRIBUTORS

Victor Adeyanju · viadeyan@microsoft.com

WHAT'S INSIDE

Five themes, 18 pages

Follow the themes in order for a full rollout, or select any page to jump straight to it.

01 Orient The map, the big picture and how agents are discovered	PAGE 2 Technical Advisory Map Advisory areas, headline recommendations and where to go deeper	PAGE 3 Enterprise Reference Architecture One governed inventory across Entra, Purview and Defender	PAGE 4 Agent 365 Discovery Paths How agents reach the registry, and Shadow AI detection	
02 Govern & secure Identity, access, MCP, risk and blast radius	PAGE 5 Identity Architecture & Lifecycle Blueprints per environment, naming, sponsors and owners, lifecycle	PAGE 6 AuthN / AuthZ & MCP Governance Permissions, credentials and MCP server approval	PAGE 7 Risk, Observability & Human Oversight Agent risk, tracing, audit and oversight tiers	PAGE 8 Security Architecture & Blast Radius Layered controls and the containment ladder
03 Agent types & build How controls differ by platform, and how developers onboard pro-code agents	PAGE 9 Agent Types · Identity & Access Identity, environments and access by agent type	PAGE 10 Agent Types · Risk & Operations Risk, observability and governance by agent type	PAGE 11 Developer Guide · Pro-Code Agents Agent 365 CLI, SDK and Skills from setup to publish	
04 Configure & operate Admin center control plane, tenant rollout, connected platforms and day-2 ownership	PAGE 12 Agent 365 Control Plane Admin center areas, management rules, lifecycle actions, monitoring, Purview scope	PAGE 13 Configuration Blueprint Phased tenant configuration with exit criteria	PAGE 14 Connected Platforms Setup Least-privilege sync from Bedrock, Vertex AI, Agentforce and more	PAGE 15 Operating Model Roles, RACI and governance cadence
05 Measure & learn Prove value, learn from real rollouts, find every source and see what changed	PAGE 16 Analytics · Usage, Adoption & Impact Usage, adoption, impact and cost for admins and leaders	PAGE 17 Case Studies & Lessons Learned How Microsoft and customers are rolling out agents	PAGE 18 References Every cited source, grouped by topic, with clickable links	PAGE 19 Change Log Versions from the original diagram onward: what changed, who updated it, whose feedback

Microsoft Agent 365 Technical Advisory

Each requested advisory area, the headline recommendations, and the page in this file that covers it in depth.

Advisory area	Customer asks	Headline recommendations	Deep-dive page
Agent Identity Architecture & Blueprint Identity Governance & Lifecycle	- Multiple sponsor, owner and manager accountability model - Blueprint, blueprint principal and agent identity for a single tenant with multiple subscriptions - Identity lifecycle: draft → publish → retirement - Agent ownership transfer process - Sponsor or owner leaver / mover workflow	- One blueprint per workload per environment (trust boundary); one agent identity per logical agent - Sponsor = accountable business group + named person; at least 2 owners; manager only on an agent's user account - Blueprint credential = the subscription's managed identity (federated credential); Azure RBAC goes to the agent identity - Sponsorship auto-transfers to the manager on leave; Lifecycle Workflows notify manager and cosponsors - Retire with block → soft delete (30 days) → permanent delete; deleting a blueprint cascades	Identity Architecture & Lifecycle p. 5
Authentication & Authorization Credential Management	Best practices for authentication and authorization	- Autonomous = client credentials + application permissions; interactive = on-behalf-of + delegated - Prefer delegated over application permissions; use Azure RBAC at resource scope - Managed identity (federated) → certificate in Key Vault → client secret for dev only - Microsoft Entra blocks high-privilege roles and Graph permissions for agents; don't work around it with app registrations	AuthN / AuthZ & MCP Governance p. 6
MCP Governance	MCP registration, governance and access control	- Register BYO MCP servers with the Agent 365 CLI; prefer EntraOAuth - Evaluate, then admin approves in Agent Tools and grants Entra consent (AI Administrator) - Calls route through the Agent 365 Tooling Gateway; block at server level at runtime - Hunt <i>ExecuteToolByGateway</i> in Defender; discover unapproved MCP configs on endpoints	AuthN / AuthZ & MCP Governance p. 6
ID Protection and Conditional Access & Risk	Best practices for agent risk management	- Baseline: a Conditional Access policy that blocks agents unless approved - Block high agent risk for all agent identities (template available) - Triage the Risky agents report: confirm compromise / safe, dismiss, disable - Correlate registry risk signals from Entra, Purview and Defender	Risk, Observability & Human Oversight p. 7
Observability & Audit	- Agent decision traceability - Agent activity tracking - Shadow AI discovery - Monitor all agents without identity	- Trace runs with OpenTelemetry spans: invoke_agent → chat → execute_tool - Entra sign-in / audit logs + Purview audit for every agent interaction type - Shadow AI page + Defender local AI agent inventory - Reconcile Registry <i>Unmanaged agents</i>, connected platforms and Defender inventory against Entra agent identities	Risk, Observability & Human Oversight p. 7 Enterprise Reference Architecture p. 3
Governance & Policy Enforcement	- Organizational policy enforcement - Global governance policy model - Human-in-the-loop (HITL) - Human-on-the-loop (HOTL) - Human-out-of-the-loop governance	- Global model: tenant settings + policy templates + management rules + Entra CA, applied by agent risk tier - Tier every agent by autonomy and impact; HITL for irreversible or external actions (tool approval) - HOTL = act autonomously with live monitoring and a tested kill switch - Human-out-of-the-loop only for low-impact, reversible, tightly scoped agents with sponsor attestation	Risk, Observability & Human Oversight p. 7 Agent 365 Control Plane p. 12 Configuration Blueprint p. 13
Security Architecture	Security architecture from agent identity and blast radius control	- Layered controls: identity, permissions, tools, data, network, runtime, response - Containment ladder from narrowest to broadest: tool → agent → blueprint → attribute-scoped CA → all agents - Protect blueprint credentials: Conditional Access doesn't apply when a blueprint creates identities - Separate Dev / Test / Prod blueprints and managed identities; Environment attribute + CA fence prod	Security Architecture & Blast Radius p. 8

Agent-type differences (Agent Builder, Copilot Studio, Foundry, custom pro-code) for every area: see the two **Agent Types** pages (p. 9–10). Recommendations are grounded in Microsoft Learn. HITL / HOTL / human-out-of-the-loop is an advisory tiering model mapped to Microsoft controls, not a product feature. Several capabilities are preview or Frontier; validate in the customer's tenant.

IN THIS DOCUMENT — SELECT A PAGE TO OPEN IT

2 Technical Advisory Map	3 Enterprise Reference Architecture	4 Agent 365 Discovery Paths	5 Identity Architecture & Lifecycle	6 AuthN / AuthZ & MCP Governance	7 Risk, Observability & Human Oversight
8 Security Architecture & Blast Radius	9 Agent Types · Identity & Access	10 Agent Types · Risk & Operations	11 Developer Guide · Pro-Code Agents	12 Agent 365 Control Plane	13 Configuration Blueprint
14 Connected Platforms Setup	15 Operating Model	16 Analytics · Usage, Adoption & Impact	17 Case Studies & Lessons Learned	18 References	19 Change Log

Microsoft Agent 365 — Enterprise Reference Architecture & Best Practices


Every onboarding route converges on one governed inventory. Numbered badges map to the best-practice panel on the right.




Microsoft Agent 365 — Agent Discovery Paths

Registry ingestion, inventory access, and Shadow AI detection are separate workflows

AGENT 365 REGISTRY INGESTION



SHADOW AI DETECTION (FRONTIER PREVIEW) — SEPARATE EXPERIENCE, NOT REGISTRY INGESTION



Scope notes: Shadow AI is Frontier preview · scheduled connected-platform sync is a future release · Graph APIs and Agent Map read existing inventory, they are not ingestion · Best practices: see Enterprise Reference Architecture (p. 3)

LEGEND

→ registry ingestion

— — ➔ read-only inventory access

→ Shadow AI detection

Microsoft-native source

Third-party / partner source

Developer / manual registration

Path / mechanism step

Agent 365 control plane

Read-only inventory view

Agent Identity Architecture, Accountability & Lifecycle

Microsoft Entra Agent ID object model, how to lay out Dev, Test and Prod across tenants and subscriptions, who is accountable, and how identities move from draft to retirement.

A · OBJECT MODEL

Microsoft Entra tenant

Agent identity blueprint (template)

- Holds the credentials: federated (managed identity), certificate or secret
- Inheritable permissions; Conditional Access and policies apply to every child
- Metadata: name, publisher, description, tags
- Requires a sponsor; set owners
- Disable it = kill switch for all its identities

added to tenant

creates · issues tokens

Agent identity (one per logical agent)

- No credentials of its own; tokens are issued through the blueprint
- Own sponsor, owners, custom security attributes
- Distinct sign-in and audit entries
- Gets Azure RBAC, Graph permissions, access packages

optional pairing

Blueprint principal (per tenant)

- Created when the blueprint is added to a tenant
- The token *oid* and the audit-log actor when the blueprint creates identities
- Single-tenant: blueprint and identities live in the same tenant
- Soft delete cascades to child identities

Agent's user account (optional, 1:1)

- Only when a user object is required: mailbox, Teams, calendar
- Has a UPN and a **manager**; may need licenses
- Separate Conditional Access principal type
- Use the same sponsor as its agent identity

Naming and tagging convention (suggested; keep it identical in every environment)

Object	Pattern	Example
Blueprint	BP-<Workload>-<Env>	BP-Finance-Prod · BP-Finance-Test · BP-Finance-Dev
Agent identity	<workload>-<agent>-<env>	finance-invoice-agent-prod
Managed identity (credential)	MI-<Workload>-<Env>	MI-Finance-Test (one per environment)
Sponsor group	SG-Agent-<Workload>-Sponsors	Same business group across environments
Custom security attributes	Environment · Department · DataSensitivity	Environment = Dev Test Prod
Registry display name	Suffix non-prod agents	Invoice Agent [TEST] — users and admins can tell them apart

Attribute-based Conditional Access (Learn: “attribute-driven Conditional Access”). Scoping Conditional Access to custom security attributes lets security teams enforce **global guardrails**[32][33][36][26], for example automatically blocking any agent tagged **Environment = Dev** or **Environment = Test** from getting tokens for production resources: Entra-protected APIs, MCP servers and databases, plus SharePoint Online and Microsoft Graph as whole services (not per site). Newly tagged agents are covered automatically. Not the same as Azure ABAC role-assignment conditions.

Suggested policy set (suggested names; Options A and B; string-type attributes only; report-only first)

Policy	Assign to	Target resources	Control
CA-Agents-Dev-Isolation	Agent identities with Environment = Dev, or the BP-*->Dev blueprints	All resources; exclude resources tagged Environment = Dev	Block
CA-Agents-Test-Isolation	Agent identities with Environment = Test	All resources; exclude Test-tagged and approved shared resources	Block
CA-Agents-Block-HighRisk	All agent identities	All resources · condition: agent risk = High	Block
CA-Users-Agent-Testers	Test accounts (e.g. SG-Agent-Testers) for on-behalf-of testing	Resources the agents call	MFA strength; agent policies don't apply to OBO

B · ENVIRONMENTS · TENANTS · BLUEPRINTS · SUBSCRIPTIONS

One blueprint per workload per environment. Tenants are identity boundaries; subscriptions are resource boundaries.

OPTION A · DEV / SANDBOX TENANT

DEV

BP-Finance-Dev · BP-Shared-Dev

- Custom security attribute: Environment = Dev
- Test users and synthetic data only
- Experiment with auth flows, CA and consent
- No access to prod resources

promote

federated credential

Sub: Dev

Sandboxes · developer machines
MI-<Workload>-Dev

PRODUCTION MICROSOFT ENTRA TENANT

TEST

BP-Finance-Test · BP-Shared-Test

- Custom security attribute: Environment = Test
- Prod-like CA, templates and consent
- Masked or synthetic data
- Pilot audience in the registry

promote

federated credential

Sub: Test

CI/CD pipelines · pre-prod hosting
MI-<Workload>-Test

PROD

BP-Finance-Prod · BP-Shared-Prod

- Custom security attribute: Environment = Prod
- Business sponsor group + named person
- Only the pipeline deploys; no human credentials
- Split further by sensitivity

federated credential

Sub: Prod-Finance · Prod-Shared

Foundry · App Service · Functions
MI-<Workload>-Prod

Promote code and configuration, never identities or credentials. Each environment gets its own blueprint, identities, consent, Conditional Access and Azure RBAC, created by the pipeline from the same source-controlled definition.

Only one tenant? Option B: skip the Dev tenant and run Dev in the production tenant (supported; compensating controls replace the tenant boundary)

- Create **BP-<Workload>-Dev** beside Test and Prod (minimum two: NonProd and Prod). The pipeline tags the blueprint *and* every agent identity **Environment = Dev**
- **Block Dev agents** (Environment = Dev, or the Dev blueprints) from **All resources**, excluding only resources tagged Environment = Dev (dev APIs, MCP servers). This also blocks SharePoint, Exchange and Microsoft Graph (Graph can't be targeted on its own)
- Conditional Access decides per resource (API), not per site or mailbox. If Dev agents need Microsoft 365 data, exclude that resource and limit permissions to test content (e.g. Sites.Selected on test sites)
- It applies when agents use **their own identity**. On-behalf-of calls run with the user's access, so test with dedicated test accounts, not admins. API keys bypass Conditional Access
- Block is the only control; start in report-only. Only Attribute Assignment Administrators set tags, so developers can't relabel Dev as Prod
- Copilot Studio shares one Microsoft-owned blueprint: tag each agent identity, and use Power Platform environments and data policies

[Learn: Conditional Access for on-behalf-of agents ↗](#)

[Learn: Filter for applications \(resource tags\) ↗](#)

Best practices for environments

- **Prod always has its own blueprints.** Never share a blueprint or its credential between prod and non-prod
- **Dev and Test get separate blueprints** when they run on different runtimes, secrets or networks (the usual case). Use one *BP-<Workload>-NonProd* only if they really share a trust boundary
- Split further by data sensitivity (e.g. Finance vs Shared), not by scale-out, audit or memory needs
- Prefer a **separate dev tenant** (Option A) when app-only admin consent, CA experiments or test data mustn't touch prod; keep Test in the prod tenant for realistic pre-prod. With one tenant, use Option B
- Scope Conditional Access by the **Environment** attribute so Dev/Test identities can't reach prod resources; grant Azure RBAC to each environment's **agent identity** at resource scope
- Non-prod agents in the prod tenant appear in the Agent 365 registry and count against directory quota: label, tag and clean them up
- Foundry: unpublished agents share the project identity, so use a separate project per environment; publishing creates that project's own blueprint
- Keep blueprint definitions, permissions and each environment's a365 config in source control; rebuild, don't hand-edit
- Define the attribute taxonomy once (Environment, Department, DataSensitivity) and tag every agent at creation; untagged agents escape attribute-based policies, so report on them weekly

C · MULTIPLE SPONSOR, OWNER AND MANAGER ACCOUNTABILITY MODEL

Role	Applies to	Can do	Can't do	Limits and rules	Best practice
Sponsor (business)	Blueprint, blueprint principal, agent identity, agent's user account	Decide renew / extend / retire; request access packages; approve in reviews; disable or soft-delete	Can't change app settings; can't re-enable or restore	Up to 100 sponsors, max 5 groups. Groups: dynamic (security or Microsoft 365) or assigned Microsoft 365; not role-assignable or assigned security. Required at creation.	Use a business-team group plus a named accountable person; keep the same sponsor on the agent identity and its user account
Owner (technical)	Blueprint, blueprint principal, agent identity	Configure, operate, respond to incidents; manage credentials	No business-lifecycle authority by default	Set at creation (best practice)	At least 2 owners from the engineering team; you can't remove the last owner
Manager (org hierarchy)	Agent's user account only	Sees agents reporting to them; requests access packages for the agent's user account	Can't modify or delete the agent	One per agent's user account	Set it for agents with a mailbox or Teams presence
Admin roles	Tenant	Agent ID Administrator / Developer create and manage; AI Administrator governs in the Microsoft 365 admin center	—	Least privilege; eligible, not permanent	Separate duties between builders (Developer) and governors (Administrator)

D · IDENTITY LIFECYCLE: DRAFT → PUBLISH → OPERATE → TRANSFER → RETIRE

DRAFT

- Builder uses Agent ID Developer role
- Dev blueprint (BP-<Workload>-Dev) in the dev tenant (A) or tagged Environment = Dev in the prod tenant (B)
- Foundry: shared project identity in the Dev project
- Sponsor and owners set at creation

PUBLISH

- Pipeline promotes to Test, then Prod; Prod identities are created fresh, never copied
- Admin approval in Requests; policy template and pilot audience
- Production handshake: sponsor, permissions, Conditional Access, groups verified

OPERATE

- Time-bound access packages
- Access reviews: sponsor attests every 6–12 months
- Quarterly review: orphaned or stale identities, and any Dev/Test identity with prod access
- Rotate certificates at least yearly

TRANSFER

- Leaver: sponsorship auto-moves to the sponsor's manager
- Lifecycle Workflows (mover / leaver): email manager and cosponsors
- Owners: add the new owner first, then remove the old one
- Agent Builder / Copilot Studio: **Assign new owner** or the reassign-ownerless rule

RETIRE

- Block in the Microsoft 365 admin center, or disable the identity
- Delete → soft delete, restorable for 30 days
- Deleting a blueprint cascades to its identities and user accounts
- Retire the Dev/Test blueprints too; update the CMDB

Ownership transfer and leaver / mover workflow



Microsoft 365 admin center owner reassignment applies to shared Agent Builder and Copilot Studio agents; the previous owner loses all access. Copilot Studio deleting an agent removes its Agent ID. Entra ID Governance for agents is in preview.

Authentication, Authorization, Credentials & MCP Governance

How agents prove who they are, what they're allowed to do, and how MCP tools are registered, approved and controlled.

A · AUTHENTICATION PATTERNS

Pattern	OAuth flow	Permissions	Consent	Token subject	Use for
Autonomous	Client credentials	Application	Admin	Agent identity	Background jobs, system-to-system
Interactive	On-behalf-of (OBO)	Delegated	User or admin	User, with the agent as actor	Chat assistants on user data
Agent's user account	Agentic user	User-style	Admin	Agent's user account	Needs a mailbox or Teams presence

B · CREDENTIAL LADDER

PREFER

Managed identity as a federated credential on the blueprint — no stored secret, Azure rotates it

ACCEPT

Certificate in Azure Key Vault or an HSM; rotate at least yearly

DEV ONLY

Client secret — replace before production

Credential rules

- Isolate credentials per blueprint and per environment: MI-<Workload>-Dev / -Test / -Prod
- Alert on expiry and on credentials added outside the pipeline
- Check sign-in logs for the credential type actually used

Authentication and authorization best practices

- Always use Agent ID (blueprint → agent identity); never plain app registrations, service principals or human user accounts for agents
- Prefer delegated over application permissions; avoid tenant-wide write
- Agent-specific Conditional Access (agents can't do MFA); use attribute-based policies on custom security attributes as global guardrails (p. 5)
- Use the Entra Auth SDK sidecar or Microsoft.Identity.Web.AgentIdentities rather than hand-rolled token code
- Set the correct token **audience** for each downstream service (for example https://storage.azure.com)
- Audit consented permissions on each blueprint quarterly to prevent privilege creep

C · AUTHORIZATION DECISION GUIDE

What does the agent need?

- Azure resource** → Azure RBAC on that resource or resource group, assigned to the agent identity
- A user's data** → delegated Graph permissions (OBO); the agent can't exceed the user
- Tenant-wide data, no user** → application permissions, sparingly, admin-consented
- Directory action** → only an allowed lower-privilege Entra role
- Common access for a fleet** → access packages with expiry
- Every instance of a class** → inheritable permissions on the blueprint

Blocked for agents: Global Administrator, Privileged Role Administrator, User Administrator, custom roles, role-assignable groups, Application.ReadWrite.All, RoleManagement.ReadWrite.All, User.ReadWrite.All, Directory.AccessAsUser.All

D · MCP REGISTRATION & GOVERNANCE FLOW

1 · **Register** — Developer runs *a365 develop-mcp register-external-mcp-server*: URL, auth type, declared tools

2 · **Evaluate** — CLI evaluation scores tool names, descriptions and schemas; fix before submitting

3 · **Review** — Microsoft 365 admin center › Agent Tools: review the server and its **tools snapshot**

4 · **Approve + consent** — AI Administrator (or Global Admin) approves and grants the Entra permissions; up to 30 min to appear

5 · **Use via gateway** — Calls route through the Agent 365 Tooling Gateway: Copilot Studio, VS Code, Claude Code, GitHub Copilot CLI

6 · **Monitor** — Defender advanced hunting: CloudAppEvents where ActionType = ExecuteToolByGateway

7 · **Block / retire** — Server-level block is enforced at runtime on every client; delete when retired

E · MCP AUTH TYPE PREFERENCE

EntraOAuth

Preferred — Entra tokens, Conditional Access, agent identity

ExternalOAuth

When the server uses a non-Entra identity provider

APIKey

Avoid; vault the key and rotate it

NoAuth

Public, read-only data only

Access control

- Approve per server and read the tools snapshot
- Expose the fewest tools; scope each tool narrowly (no “run any query” tools)
- Tool approval (human-in-the-loop) on irreversible or external actions
- Block plugins and skills you don't need


MCP governance best practices

- Maintain an MCP allow-list: only gateway-registered, approved servers in production
- Discover unapproved local MCP configurations on endpoints with Defender for Endpoint
- For self-hosted servers, front them with Azure API Management and follow Microsoft's MCP security guidance
- Review third-party MCP terms and data flows; you're responsible for data leaving your boundary

BYO MCP servers and the Agent 365 Tooling Gateway are in preview. Supported clients and auth types can change; confirm in Microsoft Learn before rollout.

Agent Risk, Observability, Audit & Human Oversight

Detect and contain risky agents, trace every decision, find agents without an identity, and match human oversight to agent autonomy.



Security Architecture from Agent Identity & Blast Radius Control

Defense in depth anchored on the agent identity, and a rehearsed containment ladder from the narrowest to the broadest kill switch.

A · DEFENSE-IN-DEPTH LAYERS (ANCHORED ON AGENT IDENTITY)

IDENTITY

Agent ID for every agent · one identity per logical agent · blueprint per trust boundary · sponsor + owners · custom security attributes (Environment, Department, DataSensitivity)

ACCESS

Baseline Conditional Access block-unless-approved · block high agent risk · delegated over application permissions · Azure RBAC at resource scope · time-bound access packages

PRIVILEGE

Platform blocks high-privilege roles and Graph permissions for agents · no custom roles or role-assignable groups · inheritable permissions reviewed quarterly

TOOLS / MCP

Approved MCP servers only, through the Tooling Gateway · minimal, narrowly scoped tools · human approval on irreversible actions · server-level block

DATA

Sensitivity labels (agents need VIEW + EXTRACT) · restrict content discovery · DLP and Insider Risk scoped to agent instances · retention

NETWORK

Global Secure Access network visibility · compliant network for agent user accounts · Endpoint DLP for gen-AI sites

RUNTIME

Defender Security for AI · Copilot Studio real-time protection · local agent runtime protection · Intune block for Shadow AI

DETECT & RESPOND

ID Protection risky agents · registry risk signals · SOC playbooks · sign-in / audit logs to SIEM · agents in incident post-mortems

B · BLAST RADIUS CONTAINMENT LADDER

Start at the narrowest step that stops the threat; move down only if needed.

1

Block one tool or MCP server

Microsoft 365 admin center > Agent Tools · *Scope: One capability, every agent that uses it*

2

Block the agent

Admin center > All agents > Block · *Scope: One agent for all users*

3

Disable the agent identity / confirm compromise

Entra > Agent ID or ID Protection · *Scope: One agent's token access*

4

Remove access (RBAC, access package, consent)

Azure portal · Entra ID Governance · *Scope: What a compromised agent can reach*

5

Disable the blueprint

Entra > Agent ID > blueprint · *Scope: Every instance of that agent class*

6

Attribute-scoped Conditional Access block

Entra > Conditional Access (custom security attribute filter) · *Scope: A segment, e.g. Environment = Dev or Test*

7

Block all agent identities (emergency)

Conditional Access > All agent identities > Block · *Scope: Every agent in the tenant*

Blast radius design principles

- **Isolate by trust boundary:** separate blueprints and credentials per environment (Dev / Test / Prod) and sensitivity; one compromise can't spread to peers
- **Unique identities:** disable one agent without affecting others; clean attribution in logs
- **Default deny:** a block-unless-approved Conditional Access policy for agent identities and agent user accounts
- **Least privilege by design:** scope to specific sites, resources and APIs; no tenant-wide write
- **Human gate on consequences:** irreversible actions require approval

Protect the control plane

- Blueprint credentials are the crown jewels: Conditional Access doesn't apply when a blueprint acquires tokens to *create* identities
- Use managed identities (federated) for blueprints; alert on new credentials and permission grants outside the pipeline
- Limit Agent ID Administrator and AI Administrator; make them eligible (PIM), not permanent
- Rehearse the ladder quarterly in non-prod; record time-to-contain
- Remember blocked SharePoint and Foundry agents are only blocked in Copilot Chat; also restrict at the source

Allowed Entra roles and blocked Graph permissions for agents evolve over time; see Authorization in Microsoft Entra Agent ID.

How the Advisory Topics Differ by Agent Type — Identity, Lifecycle, Environments, Access & MCP

Agent Builder · Copilot Studio · Microsoft Foundry · Custom pro-code (Agents Toolkit declarative agents and Agents SDK custom engine agents)

Topic	Agent Builder	Copilot Studio	Microsoft Foundry	Custom pro-code Agents Toolkit · Microsoft 365 Agents SDK
What it is · runs as	- Declarative agent built in natural language inside Microsoft 365 Copilot - Runs on the Copilot orchestrator as the signed-in user — no new privileges	- Low-code agents (declarative, custom engine, business process) in Power Platform environments - Connector calls use end-user or maker (author) connections, mediated by Power Platform	- Agents in a Foundry project (prompt, workflow, hosted) running on Foundry Agent Service in Azure - Calls downstream services as its agent identity	Declarative (Agents Toolkit): runs on the Copilot orchestrator as the user Custom engine (Agents SDK): your code and AI stack, hosted by you (Azure Bot Service)
Entra identity & blueprint	- No Entra Agent ID or app registration today - Identity = the user; governed through the registry, not Entra Agent ID	- New agents get an Entra Agent ID automatically (legacy agents on app registrations migrate automatically) - Microsoft-owned Copilot Studio blueprint (federated credential); no one in the tenant can mint tokens; you can't bring your own - Counts against directory quota; the 250-per-blueprint cap doesn't apply	- First agent in a project: default blueprint + shared project identity for all unpublished agents - Publishing (or deploying a hosted agent) creates a dedicated blueprint and identity - The project managed identity authenticates the blueprint	- Declarative: no identity of its own (user context) - Custom engine: bot app registration / managed identity by default - Add Agent ID with the Agent 365 SDK + CLI — a customer-owned blueprint you design per trust boundary
Sponsor · owner · manager	- Creator is the owner; add co-owners (you can't remove the last one) - Admin Assign new owner; rule reassigns ownerless agents to the previous owner's manager - No sponsor or manager objects (no Agent ID)	- Maker is added as sponsor on the Agent ID (owner on legacy app registrations) - Add co-authors in the environment; admin can reassign ownership - Add a business-group sponsor in Entra for continuity	- Assign sponsor and owners on the dedicated blueprint and identity in Entra after publishing - Foundry RBAC decides who can change the project - Include it in the production handshake	- You set sponsor (required) and owners when creating the blueprint and identity - Repo or pipeline owners are the technical owners - Manager only if the agent has its own user account (Frontier)
Lifecycle: draft → publish → retire	- Draft in Copilot → share (governed by the Sharing setting) - Drafts aren't visible in the registry - Retire: block or delete in the admin center (soft delete, 30 days)	- Solutions and pipelines move the agent between environments (ALM) - Drafts are visible in the registry; publish → Requests → publish with template - Deleting the agent removes its Agent ID	- Unpublished (shared identity) → published (dedicated identity) → Requests when published to Microsoft 365 - Deleting the agent application removes its identity; deleting the project removes the default blueprint	- Build in Agents Toolkit + CI/CD → upload package (Add agent or Requests) → deploy - Retire: remove from catalog, delete the blueprint (cascades), decommission Azure resources
Environments: Dev · Test · Prod	- No environments and no blueprint; a single tenant-wide agent - Test with a narrow Sharing scope or pilot group before widening - Need formal Dev / Test / Prod ALM? Build it in Copilot Studio instead	- Separate Power Platform environments (Managed Environments) with their own data policies - One Microsoft-owned blueprint; the agent in each environment has its own Agent ID (Settings › Advanced › Metadata) - Tag each Agent ID with Environment; only the pipeline deploys to Prod	- Separate Foundry projects or resources in Dev, Test and Prod subscriptions - Unpublished agents share the Dev project identity; publishing in each project creates that environment's blueprint and identity - Grant RBAC per environment; never point a Dev project at prod data	- Customer-owned BP-<Workload>-Dev / -Test / -Prod - Run a365 setup per environment with its own config and managed identity - Dev in a sandbox tenant, or in the prod tenant tagged Environment = Dev with a CA block; Test and Prod in production
Authentication & authorization	- Delegated only — can't exceed the user's access - Knowledge from SharePoint and Copilot connectors is permission-trimmed; labels honored - Main risk is oversharing: fix SharePoint permissions first	- Connector scopes attached to the Agent ID at publish (whole connector or per operation) - Advanced connector policies and data policies decide what runs - Conditional Access on the Agent ID is enforced only in Teams today	- Grant Azure RBAC to the agent identity with the right audience; the project managed identity isn't the resource principal - On-behalf-of for user-context tools - Foundry Agent Consumer role for callers	- Declarative: plugin auth (Entra SSO / OAuth) in the manifest - Custom engine: Entra SSO + OBO; bot credential: managed identity, federated or certificate - With Agent 365 SDK: service-to-service, OBO or agentic-user modes
MCP & tools	- Limited to Agent Builder capabilities, knowledge and admin-available connectors - Govern with Agent Tools blocks and Allowed agent types	- MCP servers through connectors, subject to connector and data policies - Approved BYO MCP servers from Agent Tools appear in Copilot Studio (up to 30 min) - Custom connectors, MCP and REST tools don't add scopes to the Agent ID	- MCP and A2A tools can authenticate with the agent identity (set the correct audience) - Register external servers with the Agent 365 Tooling Gateway; front self-hosted ones with API Management	- Declarative: MCP plugins use dynamic tool discovery — tools can change without republishing, so control the server - Custom engine: governed MCP tooling through the Agent 365 SDK and Tooling Gateway

Registry type mapping: Agent Builder → *Agent Builder* · Copilot Studio → *MCS DA / MCS CEA / MCS BP* · Foundry → *Foundry LOB / non-LOB / hosted* · Agents Toolkit → *Agent Toolkit* · any agent extended with the Agent 365 SDK → *Agent instance*.
Rule of thumb: if the agent has no Entra Agent ID (Agent Builder, declarative agents), govern it through the user, the data and the registry. If it has one (Copilot Studio, Foundry, SDK agents), apply the full identity controls on the other pages.

How the Advisory Topics Differ by Agent Type — Risk, Observability, Governance & Blast Radius

Agent Builder · Copilot Studio · Microsoft Foundry · Custom pro-code (Agents Toolkit declarative agents and Agents SDK custom engine agents)

Topic	Agent Builder	Copilot Studio	Microsoft Foundry	Custom pro-code Agents Toolkit · Microsoft 365 Agents SDK
Conditional Access & ID Protection	- No agent identity: the user's Conditional Access and user risk apply - Agent risk detections and agent Conditional Access don't apply	- Agent ID is covered by agent Conditional Access and ID Protection risky agents - Conditional Access on the Agent ID is enforced in Teams only today; other channels rely on connector and data policies	- Full agent Conditional Access and ID Protection on the dedicated identity - Block high risk; tag with custom security attributes (Environment, DataSensitivity)	- With Agent ID: full agent Conditional Access and ID Protection - Without it, a custom engine bot is a plain app registration — no agent risk signals; migrate with the Agent 365 SDK
Observability & audit	- Purview unified audit, usage reports and registry metadata - DLP and retention for Copilot interactions	Automatic Agent 365 telemetry: agent and tool (connector) invocations - Appears in the admin center, Defender and Purview; plus Power Platform analytics	- Foundry tracing (OpenTelemetry) - Configure Agent 365 data collection and grant observability permissions - Entra sign-in logs for the agent identity	- Declarative: inherits Copilot audit - Custom engine: instrument with the Agent 365 SDK / OTel distro (invoke_agent, chat, execute_tool) or it's invisible to Agent 365
Governance & policy enforcement	Sharing, User access, Allowed agent types - Management rules: reassign or block ownerless agents - Block / delete; pinned agents	- Power Platform: environment strategy, managed environments, data (DLP) policies, advanced connector policies, ALM - Plus admin-center Requests and policy templates	- Azure RBAC and Azure Policy on subscriptions and projects - Admin-center Requests when published to Microsoft 365; templates apply to agent identities	- App catalog and Requests; templates at activation - CI/CD gates, code review, blueprints and permissions as code
Human oversight (HITL / HOTL)	- Inherently human-in-the-loop: every turn is user-initiated and the user reviews output - Keep sensitive operations out of reach	- Workflows: Request for information human review pauses the flow and collects input (Outlook) - Agent node can request human help when unsure - Autonomous triggers = human-on-the-loop: monitor with Agent 365 and Defender	- Agent Framework tool approval (approval_mode = always_require) and workflow human-in-the-loop - Hosted agents running autonomously need human-on-the-loop monitoring and a tested kill switch	- Custom engine on Agent Framework: tool approval (ApprovalRequiredAIFunction) - Declarative: design for breach when using untrusted data; no sensitive actions without human intervention
Blast radius & kill switch	- Bounded by the user's own permissions - Kill: block in the admin center (all host apps), restrict sharing, delete - Reduce with SharePoint permission hygiene and Restricted Content Discovery	- Bounded by connector scopes and data policies per environment - Kill: block in the admin center, disable the Agent ID, tighten data policies - Isolate high-risk agents in dedicated environments	- Bounded by the identity's Azure RBAC and Graph permissions - Kill: disable the identity or blueprint, remove RBAC, delete the agent application - Admin-center block only affects Copilot Chat availability	- Bounded by your blueprint design and granted permissions - Kill: disable the blueprint (every instance), revoke consent, block or remove from the catalog, stop hosting


WHICH TYPE FITS — AND WHAT IT MEANS FOR THE CONTROL PLANE

Agent Builder	Copilot Studio	Microsoft Foundry	Custom pro-code
- Personal and team productivity on Microsoft 365 data - No new privileges; fastest path - Govern through sharing, data hygiene and the registry	- Business-process and departmental agents with connectors - Needs Power Platform environment and DLP strategy - Agent ID is automatic; add a business sponsor	- Pro-code or model-centric agents on Azure data and services - Needs subscription / project design and RBAC - Publish before production to get a dedicated identity	- Full control of stack, channels and hosting - Declarative for Copilot-hosted; custom engine for your own runtime - Add Agent 365 SDK for identity, observability, tooling

Copilot Studio Agent ID behavior reflects agents created after May 2026; legacy agents migrate automatically. Agents with their own user account are Frontier only. Validate channel-specific Conditional Access enforcement before relying on it.

Developer Guide — Bringing Custom Pro-Code Agents into Agent 365

Register, instrument, validate and publish with the Agent 365 CLI, SDK and Skills. Works for agents hosted in Azure or elsewhere.



TWO WAYS TO RUN IT — AND WHAT EACH CAPABILITY NEEDS

Agent 365 Skills in GitHub Copilot CLI, VS Code agent mode or Claude Code
`gh skill add microsoft/agent365-skills` — then describe the outcome, for example "*register this agent with Agent 365*"

Skill	What it does
a365-setup	Checks prerequisites, detects the stack, routes to the standard or AI teammate path
make-a365-agent	Registers the blueprint, agent identity and permissions
make-ai-teammate	Adds messaging and notifications; agent user account (Frontier)
instrument-observability	Adds OpenTelemetry export matched to the auth mode
add-workiq-tools	Connects Work IQ MCP servers (Mail, Calendar, Word ...)
purview-dlp-integration	Adds Purview DLP checks before prompts reach the model

Agent 365 CLI — key commands

Command	Purpose
<code>a365 setup requirements</code>	Validate roles and prerequisites; adds missing redirect URIs on a tenant-owned CLI app
<code>a365 setup all [--agent-name]</code>	Full setup (config-free with <code>--agent-name</code>); prints Global Admin next steps
<code>a365 setup blueprint</code>	Blueprint only; <code>--endpoint-only</code> / <code>--update-endpoint</code>
<code>a365 query-entra blueprint-scopes</code>	List delegated and application permissions on the blueprint
<code>a365 query-entra instance-scopes</code>	Scopes and consent status for an agent instance
<code>a365 develop-mcp register-external-mcp-server</code>	Register a BYO MCP server (CLI or <code>-f file.json</code>)
<code>a365 publish [--dry-run]</code>	Update manifest IDs and build manifest.zip for upload

Licensing by capability


Capability	Requirement
Identity	No extra license; Agent 365 enabled in the tenant
Observability	At least one user with Microsoft 365 E7 or Agent 365 assigned
Tooling (Work IQ)	Microsoft 365 Copilot license; Work IQ MCP is preview
Notifications	Agent user account; Frontier preview only

Registering existing blueprints
Microsoft 365 Agents SDK (recommended) or the CLI registers new identities in the Agent 365 registry. Blueprints created elsewhere may need adding so admins can see and govern them.

DEVELOPER BEST PRACTICES

- One blueprint per workload per environment** (Dev, Test, Prod): a separate a365 config and managed identity for each, promoted by the pipeline. Assign a sponsor and owners at creation.
 - Least privilege:** declare only the Graph scopes you need. Use OBO when acting for a user and S2S only for background work.
 - No secrets in source control.** Prefer managed identity or federated credentials; back up config files and generated IDs.
 - Emit a valid `invoke_agent` root span** or telemetry never reaches the admin center, Defender or Purview.
- Rehearse first:** use `--dry-run` before setup or publish changes, and keep the CLI current (`dotnet tool update`).
 - Don't modify the Microsoft-managed Agent 365 CLI app** (no redirect URIs or permissions); use a tenant-owned app only if needed.
 - Handle tool failures:** return a controlled error so the agent stays available; prefer EntraOAuth for MCP servers.
 - Non-Azure hosting works** (for example AWS): configure `a365.config.json` for non-Azure hosting; Entra stays the control plane.

WHO DOES WHAT



The Agent 365 Control Plane — Microsoft 365 admin center

Where admins discover, approve, automate, monitor and retire agents, and the Purview policies to scope so data protection covers them.

A · THE CONTROL PLANE AT A GLANCE — MICROSOFT 365 ADMIN CENTER › AGENTS

Overview

Agents › Overview

- 30-day snapshot: agents, active users, run-time, usage trends
- Top actions:** pending requests, agents without owners, agents at risk, agents with exceptions

Best practice: Triage top actions every week

Registry

Agents › All agents › Registry

- Every agent: Microsoft, your org, partners and connected platforms
- Filters, tags (up to 50), pinned agents, export to Excel

Best practice: Tag by workload and environment; export monthly

Requests

Agents › All agents › Requests

- Review publish requests; publish to store or reject
- Apply a policy template when you approve

Best practice: Publish to a pilot group first, then widen

Agent Map

Agents › Agent Map

- Relationships and platform clusters across your agents
- Spot unexpected connections early

Best practice: Review before widening access

Tools

Agents › Tools

- Approve BYO MCP servers and read the tools snapshot
- Block a server at runtime for every agent that uses it

Best practice: Approve per server; block what you don't need

Shadow AI

Agents › Shadow AI

- AI apps and agents detected on devices through Defender
- Block supported apps through Intune

Best practice: Detect → communicate → offer an alternative → block

Settings

Agents › Settings

- Management rules**, agent and plugin access, templates, sharing, user access, feedback sharing, tags

Best practice: Start restrictive; widen by group

B · AUTOMATED LIFECYCLE — AGENT MANAGEMENT RULES

Rules find agents that match conditions, let you review the impacted agents, then act on them in bulk (Agents › Settings › Agent management rules).

Rule	What it does	Scope and caveats	Suggested cadence
Install Microsoft agents	Bulk-installs Microsoft-published agents for all users	Review eligible agents before you run it	At rollout; after new Microsoft agents ship
Reassign ownerless agents to manager	Moves ownership to the previous owner's manager (Entra hierarchy)	Agent Builder agents only; fix Copilot Studio and other agents with Assign owner	Monthly, and after leaver waves
Block ownerless agents without usage	Blocks agents that have no owner and no usage	Review the impacted list first; blocking is reversible	Monthly
Apply a template	Applies security policies to agents that are already published	One-time run, not scheduled. Agent instances with an agent identity only; not blueprints or AI teammates	After every template change
Reject stale publish requests	Rejects requests older than the number of days you set	Tell requesters the SLA so rejections aren't a surprise	Monthly

C · LIFECYCLE ACTIONS BY STAGE — PER AGENT, FROM THE REGISTRY

Stage	Action	What it does
Onboard	Publish to store · Reject submission	Approve requested agents into the Agent Store
Onboard	Install · uninstall	Deploy to everyone or to groups; grant admin consent to the agent's permissions
Onboard	Activate	Let only selected users or groups create instances of an agent
Operate	Assign · add · remove owner	Close ownership gaps; add or remove owners on Agent Builder agents
Contain	Block · unblock	Stop everyone in the organization using the agent
Contain	Start · stop	Stop the Azure infrastructure behind a Foundry agent
Retire	Delete → restore → permanently delete	30-day recovery window, then permanent removal

Approving, installing and assigning owners needs **AI Administrator** or Global Administrator.

D · MONITORING — WHAT TO WATCH AND WHERE

Signal	Where	Action
Pending requests · agents without owners	Overview › Top actions	Approve or reject; assign an owner or run a rule
Agents at risk (Entra, Purview, Defender)	Overview · Registry <i>Risks</i> column	Investigate in the source tool; block the agent if confirmed
Exceptions · sessions · run-time	Agent details › Activity	Owner fixes errors; retire idle agents
Tool and MCP server use	Agents › Tools · Defender advanced hunting	Block servers you don't approve
Agent telemetry (OpenTelemetry)	Admin center, Defender and Purview	Developers must implement observability with the Agent 365 SDK
Central reporting	Graph API · Viva Insights advanced analysis (p. 16)	Export the registry to your CMDB; publish reports to leaders

E · PURVIEW COVERAGE FOR AGENT INSTANCES — AUTOMATIC vs. YOU MUST SCOPE IT

Capability	Coverage	What to configure
Audit · data classification · Compliance Manager AI assessments	Automatic	Search <i>Agent 365 activities</i> in Audit; review DSPM › AI observability
Data loss prevention	You must scope it	Include agent instances, or an Entra security group of them. Teams, OneDrive, SharePoint and email. Agents aren't told they were blocked, so owners monitor
Insider Risk (Risky AI usage)	You must scope it	Include agent instances as you would users
Communication Compliance · retention · eDiscovery	You must scope it	Treat agent instances like users in each policy or search
Sensitivity labels	Existing labels apply	Encrypting labels must grant the agent VIEW + EXTRACT; agent-created content isn't labeled automatically

Coverage follows the **Entra agent identity**: check each agent's Entra-enabled status in DSPM › AI observability (current DSPM; the classic version doesn't support Agent 365).

Who can do what in the admin center			
Role	View insights	View registry	Install, approve, manage
AI Administrator · Global Administrator	✓	✓	✓
AI Reader · Global Reader	✓	✓	Read-only tools and settings
Security Administrator · Security Reader · Reports Reader	✓	✓	—

Best practices

- Make **AI Administrator** the day-to-day role (eligible, not permanent); give SOC and reporting teams reader roles
- Run management rules on a fixed cadence and **review impacted agents** before each run
- Pair the admin-center block with the Entra kill switch: block stops use in Microsoft 365 hosts; disabling the agent identity stops its tokens
- Keep the Purview agent group current from the same pipeline that creates agent identities
- Reconcile the registry monthly with your CMDB, connected platforms and Defender's local agent inventory

Enterprise Configuration Blueprint — Phased Rollout

Configure in this order. Each phase lists where, who, what to configure, best practices and exit criteria (suggested).

PHASE 0 Foundations & licensing	PHASE 1 Identity foundation	PHASE 2 Registry & publishing	PHASE 3 Connect & sync platforms	PHASE 4 Protect data & threats	PHASE 5 Shadow AI	PHASE 6 Operate & improve
Where: Microsoft 365 admin center › Billing, Settings (Frontier) Who: Global Administrator (one-time), AI Administrator	Where: Microsoft Entra admin center › Agent ID, Conditional Access, ID Governance Who: Agent ID Administrator, Conditional Access / Security Administrator, Attribute Assignment Administrator	Where: Microsoft 365 admin center › Agents › Settings and All agents Who: AI Administrator	Where: Microsoft 365 admin center › Agents › All agents › Connected platforms Who: AI Administrator + cloud platform owners (AWS, Google Cloud, Salesforce, Databricks, OCI, Snowflake, Anthropic)	Where: Purview portal › DSPM · Defender portal › Settings › Security for AI Who: Compliance Administrator, Security Administrator, Power Platform Administrator	Where: Admin center › Agents › Shadow AI · Defender · Intune · Entra Global Secure Access Who: Security or AI Administrator, Intune Administrator	Where: Agent overview, Registry, Entra ID Governance, Defender, Purview Who: AI Administrator, sponsors, SOC, Compliance
CONFIGURE - Confirm Agent 365 licensing: per user, included in Microsoft 365 E7; at least one licensed user enables it - Microsoft 365 E5 is the recommended prerequisite (E5 is required to view Shadow AI) - Decide Frontier opt-in for preview features (Shadow AI, third-party observability) - Assign roles: AI Administrator (govern), AI Reader or Global Reader (audit), Security Reader (SOC) - Pick Option A (dev / sandbox tenant) or B (one tenant, Environment tags + CA block); agree the Dev → Test → Prod model	CONFIGURE - Custom security attributes: Environment (Dev Test Prod), Department, DataSensitivity - Blueprint standard: BP-<Workload>-<Env> per environment; sponsor, owners, metadata, permissions - Agent-specific Conditional Access in report-only; block high-risk agents; block Dev/Test identities from prod resources - Access packages for common agent access patterns - Access reviews and Lifecycle Workflows for sponsors	CONFIGURE Allowed agent types: Microsoft, your org, external publishers User access: specific groups for the pilot Sharing (Agent Builder): specific users or groups Templates: default plus custom (Conditional Access, access packages, custom security attributes) Tags taxonomy (max 50) and pinned agents (3 per user) Requests: Publish to store with scoped audience and template Agent Tools: approve BYO MCP servers; block unneeded tools	CONFIGURE - Create dedicated service identities (see Connected Platforms Setup page) - Connect a platform: name, region, auto-import choice, credentials, Validate, Save - Run Sync agents; review status, synced count, errors - Code-first agents: Agent 365 SDK; Foundry: configure Agent 365 data collection	CONFIGURE - Purview: confirm audit; review DSPM › AI observability - DLP, Insider Risk (Risky AI usage), Communication Compliance, retention scoped to agent instances - Sensitivity labels enabled for SharePoint and OneDrive - Defender: Microsoft 365 connector (Entra ID Management events + Microsoft 365 activities) - Copilot Studio real-time protection with the Power Platform admin	CONFIGURE - Onboard devices to Defender for Endpoint; enroll in Intune - Enable Global Secure Access traffic forwarding for usage metadata - Review detections: devices, users, first and last seen - Block supported agents (OpenClaw) → Intune policy <i>A365 - Block OpenClaw</i> - Purview Endpoint DLP for gen-AI websites	CONFIGURE - Management rules: reassign ownerless Agent Builder agents to manager; block ownerless agents without usage; reject stale requests - Monthly export or Graph pull into the CMDB - Alerts: credential expiry, Conditional Access blocks, token spikes, permission changes
BEST PRACTICES - Reserve Global Administrator for break-glass; make AI Administrator eligible, not permanent - Name an Agent 365 service owner and a governance board (IT, Security, Compliance, Business) - Publish an agent intake and acceptable-use policy before opening creation broadly - Track GA versus preview features for your tenant	BEST PRACTICES - Plan blueprints before building: one per workload per environment; one identity per logical agent - Apply Conditional Access and permissions at the blueprint; disabling it blocks every instance - Audit “all users” MFA policies for unintended agent impact - Certificates or managed identities in production; Key Vault; rotate at least yearly; separate dev / test / prod - Naming convention, for example BP-HR-Prod and hr-onboarding-agent-prod - Sponsor groups: dynamic or Microsoft 365 groups (role-assignable groups aren't supported)	BEST PRACTICES - Start restrictive, then widen by group as controls mature - Templates apply at activation; use the Apply template rule for existing agents with identities - Blocking a SharePoint or Foundry agent affects Copilot Chat only; also restrict at the source - Share agent feedback with developers for pilot agents	BEST PRACTICES - One connection per environment and region with a descriptive name, for example AWS-Prod-USEast1 - Least privilege; approve delete scopes deliberately - Sync on a runbook cadence (weekly and after releases) until scheduled sync ships - Prefer built-in integration or registry sync before adding the SDK	BEST PRACTICES - Scope policies with a security group that contains agent instances - Encrypting labels must grant agents VIEW + EXTRACT; “all users in org” isn't sufficient - Agent-created content isn't labeled automatically; plan default labeling - Owners monitor DLP blocks; agents aren't aware of them - Add agent scenarios to SOC playbooks	BEST PRACTICES - Detect → triage → communicate → offer an approved alternative → block - Blocking reaches all Intune-managed Windows devices; allow 15 min to 8 h and verify in Intune - Use Defender's local AI agent inventory for the wider list (CLI agents, IDEs, MCP servers)	BEST PRACTICES - Weekly: Requests, Agents at risk, exceptions, sync errors - Quarterly: orphaned-agent review and permission right-sizing - Every 6–12 months: sponsor attestation; decommission if not confirmed - Treat blueprints and permissions as code with a production handshake
EXIT CRITERIA - Licenses active - Roles assigned with least privilege - Governance charter and RACI signed off - Dev environment option (A or B) and model ready	EXIT CRITERIA - Attribute set published - Policies validated in report-only - High-risk agent block enforced - Review schedule created	EXIT CRITERIA - Settings baseline documented - Templates created - First pilot agent published through Requests	EXIT CRITERIA - In-scope platforms connected - First sync reconciled to each platform's inventory - Owners set for synced agents	EXIT CRITERIA - Template policies show enabled - Defender checklist Connected - DLP and Insider Risk tested with a pilot agent	EXIT CRITERIA - Detection coverage meets target - Decision recorded per detected app - Block policy verified	EXIT CRITERIA - No agent without an owner for more than 30 days - Every org agent has a template - Requests handled within SLA - High-risk agents remediated

Tip: Phases 1–2 are the foundation — configure identity and registry guardrails before connecting platforms or widening access. Phases 3–5 can run in parallel once the foundation is in place; phase 6 is steady-state.

Connected Platforms — Least-Privilege Setup

How to create the service identity Agent 365 uses for each platform, and the guardrails to apply.



Platform	What you provide	Minimum permissions (sync + management)	Extra for observability	Best practice
Amazon Bedrock (Classic, AgentCore)	AWS region; IAM access key ID and secret access key	bedrock: ListAgents, GetAgent, List/GetAgentAlias, List/GetAgentVersion, InvokeModel, InvokeAgent, InvokeInlineAgent, DeleteAgent. AgentCore: bedrock-agentcore List/Get and Delete for harnesses, runtimes, gateways, gateway targets, memory	Observability supported; review the Bedrock observability permissions	Dedicated IAM user per account and region; rotate access keys; Delete* scopes enable management actions from Agent 365, so approve them deliberately
Google Vertex AI	Region, project ID, full service account JSON key	Vertex AI Administrator, or a custom role with aiplatform.reasoningEngines.list, .get, .delete	BigQuery Data Viewer and BigQuery Job User, Monitoring Viewer, Cloud Trace User; enable APIs and trace export	Prefer the custom role over Vertex AI Administrator; one service account per project; rotate keys
Salesforce Agentforce	My Domain URL, consumer key and secret (External Client App)	OAuth scopes: chatbot_api, sfap_api, api, refresh_token / offline_access; Client Credentials Flow with a Run As user that has API access	—	Use a dedicated integration Run As user with only the access it needs
Databricks Genie	Workspace URL, service principal client ID and secret	Service principal with admin access in the workspace	—	Dedicated service principal per workspace; rotate the client secret
Anthropic Claude Managed Agents (preview)	Workspace ID, workspace-scoped API key	API key scoped to a dedicated workspace	Observability supported; review the Anthropic prerequisites	Create a new workspace (the default workspace doesn't expose an ID); APIs are beta, so pilot first
Oracle OCI Generative AI Agents	Private key, user OCID, tenancy OCID, fingerprint, region, compartment OCID	Option B (recommended): read generative-ai-agent-family; manage genai-agent and genai-agent-endpoint limited to DELETE	—	API-only service user (no console password); revoke by deleting the API key or removing the group; only active agents sync
Snowflake Cortex	Account identifier, database, schema, Snowflake user, private key	RSA key-pair authentication configured for the Snowflake user	—	Dedicated service user; follow Snowflake key-pair rotation

Common guardrails for every connection

- Name connections by platform, environment and region so sync results are easy to trace
- Keep production and non-production in separate connections and service identities
- Store every secret or private key in a vault with an owner and a rotation date
- Re-validate the connection after each credential rotation
- Review sync errors after every run; a silent failure means stale inventory
- Document the revoke path per platform before go-live

What connected-platform sync does and doesn't do

- **Does:** import agents into the registry for inventory and governance, without the Agent 365 SDK
- **Does:** support management actions that the platform APIs allow (for example delete)
- **Doesn't:** add Entra identity, Work IQ, notifications or code-level observability — use the SDK for those
- **Doesn't (yet):** run on a schedule — an admin selects **Sync agents**
- Third-party observability (Frontier) needs its own prerequisites per platform

Source: Microsoft Learn — Connected platforms in the Microsoft 365 agent registry; Third-party agent observability. Permissions change over time; confirm before provisioning.

Operating Model — Roles, RACI and Cadence

Who owns what once Agent 365 is live. R = responsible · A = accountable · C = consulted · I = informed.

Activity	AI Administrator	Entra identity admins	Security / SOC	Compliance (Purview)	Intune / endpoint	Cloud platform owners	Agent sponsor	Agent owner / developer
Licensing, roles, Frontier decision	A / R	C	C	C	I	I	I	I
Agent settings baseline and templates	A / R	R (custom policies)	C	C	—	—	I	I
Blueprint and agent identity standards	C	A / R	C	—	—	—	C	R
Promote agents Dev → Test → Prod	C	A	C	—	—	R (subscriptions, MIs)	C (prod sign-off)	R
Review and approve agent requests	A / R	—	C	C	—	—	C	R (submits)
Connected-platform credentials and sync	A / R	—	C	—	—	R	I	I
Conditional Access and ID Protection for agents	I	A / R	C	—	—	—	—	C
Purview DLP, Insider Risk, retention for agents	C	—	C	A / R	—	—	I	C
Defender Security for AI and incident response	I	C	A / R	C	—	—	C	C
Shadow AI detection and blocking	C	—	A / R	C	R	—	—	—
Ownerless agent remediation	A / R	C	—	—	—	—	R	C
Sponsor attestation and access reviews	I	A	—	—	—	—	R	C
Inventory export and CMDB reconciliation	A / R	—	I	I	—	C	—	—
MCP server / tool approval and blocking	A / R	C (consent)	C	C	—	—	I	R (registers)
Agent risk response (ID Protection, containment)	C	R	A / R	C	—	—	C	R

WEEKLY

- Approve or reject pending Requests
- Triage Agents at risk and agents with exceptions
- Run Sync agents and review sync errors
- Review new Shadow AI detections

MONTHLY

- Export registry or pull Graph packages API → CMDB
- Run ownerless-agent and stale-request rules
- Check Unmanaged agent count trend
- Tag hygiene and pinned-agent review

QUARTERLY

- Orphaned-agent review: missing sponsors, stale metadata, no activity
- Right-size agent permissions and access packages
- Remove stale Dev/Test identities; confirm none reach prod
- Review connected-platform service identities

6–12 MONTHS

- Sponsor attestation through access reviews
- Rotate certificates (at least yearly)
- Re-baseline templates, allowed agent types and user access

Suggested KPIs (Agent overview and Registry)

- Agents without owners — target 0 older than 30 days
- Unmanaged agents — trending down month over month
- Agents at risk — high-severity remediated within SLA
- Pending requests — age within SLA (for example 5 business days)
- % org-published agents with a policy template
- Connected platforms with a successful sync in the last 7 days
- Shadow AI: devices with unapproved agents, trending down

Role guidance

- **AI Administrator:** day-to-day governance (approve requests, assign owners, templates, rules, Graph APIs)
- **Global Administrator:** break-glass only
- **AI Reader / Global Reader:** read-only oversight for auditors and leadership
- **Security Admin / Reader / Operator:** risk signals, Defender, Shadow AI
- **Agent ID Administrator:** owners and sponsors on blueprints and identities
- **Attribute Assignment Administrator:** needed for custom security attribute policies in templates
- **IRM Analyst or Investigator:** open Purview Insider Risk risk-signal links

Key references (Microsoft Learn): Agent registry · Agent settings · Policy templates · Agent requests · Connected platforms · Choose an Agent 365 integration option · Best practices for Microsoft Entra Agent ID · Conditional Access for agents · Governing agent identities · Purview for Agent 365 · Defender security for AI agents · Shadow AI · Local AI agent discovery with Defender for Endpoint · Agent management roles. Linked boxes open their source article.

Measuring Agents — Usage, Adoption & Impact

Which report answers which question, who it's for, and how to turn agent activity into business impact.

Report or surface	Audience	Question it answers	Key metrics	Data, access & caveats
Agent 365 · Agent overview <i>Microsoft 365 admin center › Agents › Overview</i>	AI administrators, IT leads	What's in our agent estate, and what needs action now?	- Hero metrics: Agent registry count, Active users, Agent run-time (hours worked), Registry sync - Top actions: pending requests, agents at risk, agents without owners, agents with exceptions - Analytics: agents by creator, top build platforms, active users over time, trending agents	- Last 30 days; usage metrics start at Agent 365 license activation - Approving requests and assigning owners needs AI Administrator or Global Administrator - Foundry analytics cover Foundry V2 agents only
Agent details · Activity tab <i>Registry › select an agent</i>	Agent owners, AI administrators	Is this agent used, reliable and worth its run-time?	Active users, Sessions (a new session starts after 30 minutes of inactivity) Exceptions (sessions with errors), Agent run-time	- Default 30 days; adjustable date range - Requires a Microsoft 365 E7 or Agent 365 license
Copilot Agents usage report <i>Reports › Usage › Microsoft Copilot › Agents</i>	Microsoft 365 admins, adoption leads	Who uses which agents in Copilot Chat, Word, Excel and PowerPoint?	- Total, licensed and unlicensed active users; total active agents - Users and agents by creator type (your users, your org, Microsoft, third party) - Per-user, per-agent and per agent-user responses; last activity date	7 or 30 days; usage visible within about an hour - User names are concealed by default (reports privacy setting) - Excludes Cowork (own report) and SharePoint agents used in Teams
Copilot Credits report <i>Microsoft 365 admin center › Reports</i>	FinOps, service owners	What do agents cost, and who drives consumption?	- Credits used: total, cumulative and daily - Per user, per agent, per billing policy and per agent-user pair - Alerts when a user exceeds 2,000 credits	- Appears after pay-as-you-go billing is set up (Microsoft 365 or Power Platform admin center) 7 or 30 days
Viva Insights · Agent Dashboard (preview) <i>Viva Insights web app › Agent Dashboard</i>	Business leaders, change managers, analysts	Where is adoption growing, by organization and job function?	Agent 365 Dashboard view (all agents across Microsoft 365 apps): active agents, active users, sessions Copilot Agent Dashboard view: agent responses and Copilot Credits for agents in Copilot - Overview, adoption (by group, trends, agent snapshot) and per-agent deep dives	28-day window, up to 6 days' delay; privacy minimum group size - Needs 50+ Copilot licenses; the Agent 365 view also needs Agent 365 licensing - Auto-access for senior leaders; delegate to others; upload org data for richer filters
Viva Insights · Advanced analysis <i>analysis.insights.cloud.microsoft</i>	Central reporting team, CoE analysts	One central, customizable view of adoption and impact across the organization	- Power BI templates: Copilot Studio agents (top agents, agent-assisted hours), Copilot adoption, Copilot business impact - Custom queries sliced by uploaded HR attributes; auto-refresh when new data arrives - Publish finished reports to leaders in the Viva Insights web app	- Needs the Insights Analyst role; privacy minimum group size applies - Copilot Studio agents report: 50+ Copilot licenses and a published agent with usage - Up to 13 months of Copilot history; data lags about 6 days
Copilot Studio · Monitor and analytics	Makers, product owners	Is the agent effective, and how much time or money does it save?	- Sessions, effectiveness and engagement KPIs; up to 3 custom metrics per agent Savings calculator: time or cost saved per run or per tool, tracked over time - Copilot Agent Kit (Conversation KPIs) and Application Insights for deep telemetry	- Admins can turn off money-saving rules per environment (Power Platform admin center) - Savings are estimates; agree the baseline with the business owner
Foundry and custom pro-code telemetry	Developers, operations	Is the agent healthy and safe at run time?	- OpenTelemetry spans: invoke_agent → inference → execute_tool - Surfaces in the Microsoft 365 admin center, Defender and Purview through the Agent 365 SDK - Foundry tracing and monitoring in Azure (Application Insights)	- Telemetry is dropped unless at least one user has an E7 or Agent 365 license assigned - Runs without an invoke_agent root span only appear in Defender advanced hunting

METRICS FRAMEWORK — FROM ACTIVITY TO VALUE

USAGE · is it used?	ADOPTION · is it spreading?	IMPACT · is it worth it?	COST · is it efficient?
- Active users and active agents - Sessions and agent run-time - Exception rate (exceptions ÷ sessions) - Responses per user	% of target population active, by org and function (Viva) - Repeat use: sessions per active user - Mix of creator types (org-published vs user-created) - Licensed vs unlicensed use	- Time and cost saved per run or per tool (Copilot Studio) - Business KPI moved: cycle time, backlog, CSAT, quality - Agents promoted from pilot to business-critical - Sponsor-attested outcomes each quarter	- Copilot Credits per agent and per billing policy - Cost per session and per hour saved - Idle agents: run-time or credits with low usage - Spend alerts reviewed weekly

BEST PRACTICES Assign licenses first. Usage metrics start at Agent 365 license activation, and SDK telemetry is dropped until at least one E7 or Agent 365 license is assigned. Baseline before launch. Record the current time and cost of the task so savings and ROI hold up. Two or three KPIs per high-impact agent, agreed with its sponsor and reviewed monthly. Match the report to the audience: admin center to act, Viva Insights for leaders, advanced analysis for central reporting, Copilot Studio and Foundry for makers.	Turn insight into action: low usage + no owner = retire; high usage + no owner = reassign; trending agents = risk review, then scale. Segment the estate: separate short-lived experiments from business-critical agents so trends aren't distorted. Respect privacy: keep usage-report user details concealed unless there's a need; Viva applies minimum group sizes. Reconcile monthly: registry and analytics come from different pipelines, so expect small timing variances.
--	--

Case Studies — How Organizations Roll Out Agent 365 and Govern Agents

What Microsoft learned as Customer Zero, and rollout patterns from customers scaling agents on Microsoft platforms.

FEATURED · MICROSOFT DIGITAL, CUSTOMER ZERO FOR AGENT 365 (INSIDE TRACK)

500,000+

agents visible in one registry

3-part

admin model: AI · Agent ID · security & compliance teams

58,000

Cowork users' sessions and locations analyzed in minutes

Weekly

cross-team governance rhythm: IT, security, identity, product, business

1 · Foundation for administrators

- Split security, governance, AI admin and identity duties with explicit handoffs
- Treat Agent 365 as the coordination layer, not a replacement for platform admins
- Visibility first, then approval workflows matched to risk tolerance
- Avoid a bureaucratic choke point: partnership over centralization

2 · Registry as source of truth

- No deployment needed: validate registry accuracy and completeness first
- Every agent gets a clear owner and lifecycle state
- Compare exports with earlier inventories (Power Platform, SharePoint) to reconcile
- Use the SDK and Entra Agent ID so pro-code agents register automatically

3 · Visualize to act

- Prioritize: ownerless, duplicate, trending and strategic agents
- Unused + ownerless → retire; heavily used + ownerless → reassign
- Agent 365 tells IT what to do; Viva Insights tells adoption leaders how usage spreads
- Scale with APIs, rules and bulk actions, not manual review

4 · Secure with existing controls

- Extend Entra, Purview and Defender rather than inventing a new model
- Track build-time risk (broad permissions, weak config) and run-time risk (data exposure, prompt injection)
- Remediate in the owning tool; use Agent 365 for the shared risk view
- Expect more visibility to surface new issues, and plan for it

"My advice to every IT leader is to establish that governance rhythm before your agent count grows faster than your ability to manage it." — Garima Tiwari, Principal Product Manager, Agent 365 Customer Zero, Microsoft Digital

CUSTOMER STORIES — ROLLING OUT AND GOVERNING AGENTS AT SCALE

KPMG · AGENT 365

- Adopting Agent 365 to manage, monitor and secure agents across its global organization (June 2026)
- Copilot expanding to 276,000+ professionals
- Copilot Studio for business makers, Foundry for complex agents: "experiment at scale, then industrialize"

AGCO

- 900+ employee volunteers building agents in Copilot Studio
- AI literacy first; citizen makers paired with expert support
- Aim: safe AI use instead of ungoverned, disconnected tools

Brown Health

- AI Center of Excellence spent its first six months on governance: policies, acceptable use, responsible AI
- Then scaled to 24+ Copilot Studio agents for clinical and operational workflows

HealthEquity

- 80%+ of the workforce completed foundational Copilot training
- Monthly Copilot actions grew from ~50,000 to 220,000+ in six months
- Custom agents embedded in everyday workflows

EisnerAmper

- Foundry agent supporting 18,000 audits, secure by design
- Purview classifies and labels sensitive data; security groups scope access
- Security was the guiding principle from day one

PATTERNS THAT REPEAT — AND HOW TO APPLY THEM WITH AGENT 365

Pattern	Seen at	How to apply it with Agent 365	Where in this file
Visibility before control	Microsoft Digital	Validate the registry and connect platforms, then add Requests, templates and management rules	Reference Architecture p. 3 , Discovery Paths p. 4
Governance before scale	Brown Health	Define acceptable use and risk tiers; set Allowed agent types and sharing before broad enablement	Configuration Blueprint p. 13
Citizen makers with guardrails	AGCO, KPMG	Scope Agent Builder and Copilot Studio sharing to groups; pilot groups; makers become sponsors	Agent Types p. 9
Enablement drives adoption	HealthEquity	Use the Viva Insights Agent Dashboard to target training where adoption lags	Analytics p. 16
Secure by design for pro-code	EisnerAmper, KPMG	Entra Agent ID per trust boundary, Purview labels and DLP, Defender for AI	Identity p. 5 , Developer Guide p. 11
A shared operating rhythm	Microsoft Digital	Weekly cross-team review of risky, ownerless and trending agents	Operating Model p. 15

Sources: Microsoft Inside Track, Implementing Agent 365: How we're governing and managing AI agents at Microsoft; Microsoft News, KPMG and Microsoft scale trusted, enterprise AI agents globally (June 9, 2026); Microsoft Customer Stories (AGCO, Brown Health, HealthEquity, EisnerAmper, KPMG). Only Microsoft Digital and KPMG explicitly reference Agent 365; the other stories show agent rollout patterns that Agent 365 helps govern. Results are as reported by each organization. Select any card to open its source.

Microsoft Agent 365 Technical Advisory · Microsoft FastTrack · September 2026 · Based on Microsoft Learn; preview and Frontier capabilities may change

Questions or feedback: Carlos Suarez (Solution Engineer) · Alejandro Lopez (FastTrack Engineer) | Page 17 of 19

References & Sources

Every source used in this document, grouped by topic. Select a title to open it; links stay clickable in draw.io PDF exports.

Microsoft Agent 365 & the Microsoft 365 admin center · 17

[1] [Microsoft Agent 365 overview](#)

Microsoft Learn · Cited on p. 4, 13
learn.microsoft.com/microsoft-agent-365/overview

[2] [Agent overview in Microsoft 365 admin center](#)

Microsoft Learn · Cited on p. 3, 12, 13, 15, 16
learn.microsoft.com/microsoft-365/admin/manage/agent-365-overview

[3] [Agent Registry in Microsoft 365 admin center](#)

Microsoft Learn · Cited on p. 3, 4, 7, 12
learn.microsoft.com/microsoft-365/admin/manage/agent-registry

[4] [Understand agent details in Microsoft 365 admin center](#)

Microsoft Learn · Cited on p. 16
learn.microsoft.com/microsoft-365/admin/manage/agent-details

[5] [Use Agent Map in the Microsoft 365 admin center](#)

Microsoft Learn · Cited on p. 4, 12
learn.microsoft.com/microsoft-365/admin/manage/agent-map

[6] [Agent Requests in Microsoft 365 admin center](#)

Microsoft Learn · Cited on p. 3, 12
learn.microsoft.com/microsoft-365/admin/manage/agent-requests

[7] [Agent settings in Microsoft 365 admin center](#)

Microsoft Learn · Cited on p. 3, 12, 13
learn.microsoft.com/microsoft-365/admin/manage/agent-settings

[8] [Governance and lifecycle actions for agents in Microsoft 365 admin center](#)

Microsoft Learn · Cited on p. 5, 12
learn.microsoft.com/microsoft-365/admin/manage/agent-actions

[9] [Agent management roles and permissions in Microsoft 365 admin center](#)

Microsoft Learn · Cited on p. 12, 15
learn.microsoft.com/microsoft-365/admin/manage/agent-roles-perms

[10] [Manage tools for agents in Microsoft 365 admin center](#)

Microsoft Learn · Cited on p. 6, 8, 12
learn.microsoft.com/microsoft-365/admin/manage/manage-tools-for-agent

[11] [Understand Shadow AI in Microsoft 365 admin center](#)

Microsoft Learn · Cited on p. 3, 4, 12, 13
learn.microsoft.com/microsoft-365/admin/manage/agent-shadow-ai

[12] [Policy templates](#)

Microsoft Learn
learn.microsoft.com/microsoft-agent-365/admin/policy-template

[13] [Connected platforms in Microsoft Agent 365](#)

Microsoft Learn · Cited on p. 3, 4, 13, 14
learn.microsoft.com/microsoft-agent-365/admin/connected-platforms

[14] [Ecosystem partner agents available in Agent 365](#)

Microsoft Learn · Cited on p. 3, 4
learn.microsoft.com/microsoft-agent-365/third-party-agents

[15] [Third-party agent observability with Microsoft Agent 365 \(Frontier\)](#)

Microsoft Learn · Cited on p. 14
learn.microsoft.com/microsoft-agent-365/admin/third-party-agent-observability

[16] [Graph API for agent registry and agent details](#)

Microsoft Learn · Cited on p. 3, 4
learn.microsoft.com/microsoft-agent-365/admin/graph-api

[17] [Observability \(Microsoft Agent 365 admin\)](#)

Microsoft Learn · Cited on p. 12
learn.microsoft.com/microsoft-agent-365/admin/monitor-agents

Security, compliance & human oversight · 5

[18] [Use Microsoft Purview to manage data security & compliance for Microsoft Agent 365](#)

Microsoft Learn · Cited on p. 3, 8, 12, 13
learn.microsoft.com/purview/ai-agent-365

[19] [Learn about Endpoint data loss prevention](#)

Microsoft Learn · Cited on p. 3
learn.microsoft.com/purview/endpoint-dlp-learn-about

[20] [Enable security for AI agents using Microsoft Defender](#)

Microsoft Learn · Cited on p. 3
learn.microsoft.com/defender-xdr/security-for-ai/get-started-defender-security-for-ai

[21] [Local AI agent discovery with Microsoft Defender for Endpoint](#)

Microsoft Learn · Cited on p. 3, 6, 7, 8
learn.microsoft.com/defender-endpoint/local-agent-discovery-overview

[22] [Using function tools with human in the loop approvals \(Agent Framework\)](#)

Microsoft Learn · Cited on p. 7
learn.microsoft.com/agent-framework/agents/tools/tool-approval

Identity — Microsoft Entra Agent ID · 19

[23] [Fundamental concepts in Microsoft Entra Agent ID](#)

Microsoft Learn · Cited on p. 5
learn.microsoft.com/entra/agent-id/identity-platform/key-concepts

[24] [Plan your agent identity architecture](#)

Microsoft Learn · Cited on p. 5, 6, 7, 8
learn.microsoft.com/entra/agent-id/how-to-plan-agent-identity-architecture

[25] [Administrative relationships in Microsoft Entra Agent ID \(owners, sponsors and managers\)](#)

Microsoft Learn · Cited on p. 5
learn.microsoft.com/entra/agent-id/agent-owners-sponsors-managers

[26] [Best practices for Microsoft Entra Agent ID](#)

Microsoft Learn · Cited on p. 3, 5, 6, 7, 8, 13
learn.microsoft.com/entra/agent-id/best-practices-agent-id

[27] [Authorization in Microsoft Entra Agent ID](#)

Microsoft Learn · Cited on p. 6, 8
learn.microsoft.com/entra/agent-id/authorization-agent-id

[28] [Create an agent identity blueprint](#)

Microsoft Learn · Cited on p. 4, 11
learn.microsoft.com/entra/agent-id/create-blueprint

[29] [Delete and restore agent identity objects](#)

Microsoft Learn · Cited on p. 5
learn.microsoft.com/entra/agent-id/howto-delete-agent-identity

[30] [Governing agent identities \(Microsoft Entra ID Governance\)](#)

Microsoft Learn · Cited on p. 5
learn.microsoft.com/entra/id-governance/agent-id-governance-overview

[31] [Agent identity sponsor tasks in Lifecycle Workflows](#)

Microsoft Learn · Cited on p. 5
learn.microsoft.com/entra/id-governance/agent-sponsor-tasks

[32] [Conditional Access for agents in Microsoft Entra](#)

Microsoft Learn · Cited on p. 3, 8
learn.microsoft.com/entra/identity/conditional-access/agent-id

[33] [Recommended policies for autonomous agents in Microsoft Entra](#)

Microsoft Learn · Cited on p. 5, 7
learn.microsoft.com/entra/identity/conditional-access/policy-autonomous-agents

[34] [Target agent identities in Conditional Access policies](#)

Microsoft Learn · Cited on p. 5
learn.microsoft.com/entra/identity/conditional-access/howto-target-agent-identities

[35] [Conditional Access for agents operating on-behalf-of a user in Microsoft Entra](#)

Microsoft Learn · Cited on p. 5
learn.microsoft.com/entra/identity/conditional-access/policy-on-behalf-of-agents

[36] [Conditional Access: Filter for applications](#)

Microsoft Learn · Cited on p. 5
learn.microsoft.com/entra/identity/conditional-access/concept-filter-for-applications

[37] [ID Protection for agents](#)

Microsoft Learn · Cited on p. 7, 8
learn.microsoft.com/entra/id-protection/concept-risky-agents

[38] [Configure Microsoft Entra agent identities for increased security](#)

Microsoft Learn · Cited on p. 7, 8
learn.microsoft.com/entra/fundamentals/zero-trust-ai

[39] [Tutorial: Enable Internet Access traffic forwarding \(Global Secure Access\)](#)

Microsoft Learn · Cited on p. 3
learn.microsoft.com/entra/global-secure-access/tutorial-internet-access-enable-traffic-forwarding

[40] [Set up a test environment for your app \(Microsoft identity platform\)](#)

Microsoft Learn · Cited on p. 5
learn.microsoft.com/entra/identity-platform/test-setup-environment

[41] [Secure resource isolation in a single tenant in Microsoft Entra ID](#)

Microsoft Learn · Cited on p. 5
learn.microsoft.com/entra/architecture/secure-single-tenant

Agent platforms · 9

[42] [Agent Builder in Microsoft 365 Copilot](#)

Microsoft Learn · Cited on p. 9, 10
learn.microsoft.com/microsoft-365/copilot/extensibility/agent-builder

[43] [Choose between Agent Builder in Microsoft 365 Copilot and Copilot Studio](#)

Microsoft Learn · Cited on p. 9, 10
learn.microsoft.com/microsoft-365/copilot/extensibility/copilot-studio-experience

[44] [App registration, agent identities, and authentication \(Copilot Studio\)](#)

Microsoft Learn · Cited on p. 9, 10
learn.microsoft.com/microsoft-copilot-studio/govern-agents-identities-overview

[45] [Manage Entra Agent IDs \(Copilot Studio\)](#)

Microsoft Learn · Cited on p. 9, 10
learn.microsoft.com/microsoft-copilot-studio/admin-use-entra-agent-identities

[46] [Establish an Application Lifecycle Management \(ALM\) strategy \(Copilot Studio\)](#)

Microsoft Learn · Cited on p. 9, 10
learn.microsoft.com/microsoft-copilot-studio/guidance/alm

[47] [Request information from human review in workflows \(Copilot Studio\)](#)

Microsoft Learn · Cited on p. 9, 10
learn.microsoft.com/microsoft-copilot-studio/workflows-experience/flows-request-for-information

[48] [Agent identity concepts in Microsoft Foundry](#)

Microsoft Learn · Cited on p. 9, 10
learn.microsoft.com/azure/foundry/agents/concepts/agent-identity

[49] [Create and deploy a custom engine agent with Microsoft 365 Agents SDK](#)

Microsoft Learn · Cited on p. 9, 10
learn.microsoft.com/microsoft-365/copilot/extensibility/create-deploy-agents-sdk

[50] [Dynamic tool discovery for MCP plugins in Microsoft 365 Copilot](#)

Microsoft Learn · Cited on p. 9, 10
learn.microsoft.com/microsoft-365/copilot/extensibility/plugin-dynamic-tool-discovery

Developers — Agent 365 SDK, CLI & Skills · 9

[51] [Microsoft Agent 365 SDK and CLI](#)

Microsoft Learn · Cited on p. 4
learn.microsoft.com/microsoft-agent-365/developer/

[52] [Choose an Agent 365 integration option](#)

Microsoft Learn · Cited on p. 3, 4, 9, 10, 11
learn.microsoft.com/microsoft-agent-365/developer/choose-integration-option

[53] [Quickstart: Connect an existing agent to Agent 365](#)

Microsoft Learn · Cited on p. 4, 11
learn.microsoft.com/microsoft-agent-365/developer/get-started

[54] [Install and use the Agent 365 CLI](#)

Microsoft Learn · Cited on p. 11
learn.microsoft.com/microsoft-agent-365/developer/agent-365-cli

[55] [Agent 365 CLI setup command reference](#)

Microsoft Learn · Cited on p. 11
learn.microsoft.com/microsoft-agent-365/developer/reference/cli/setup

[56] [Agent 365 Skills for guided agent setup](#)

Microsoft Learn · Cited on p. 11
learn.microsoft.com/microsoft-agent-365/developer/agent-365-skills

[57] [Integrate agent observability using direct OTEL](#)

Microsoft Learn · Cited on p. 7
learn.microsoft.com/microsoft-agent-365/developer/direct-open-telemetry-integration

[58] [Agent 365 SDK validation checklist](#)

Microsoft Learn · Cited on p. 11
learn.microsoft.com/microsoft-agent-365/developer/validation-checklist

[59] [Publish agent to Microsoft admin center](#)

Microsoft Learn · Cited on p. 11
learn.microsoft.com/microsoft-agent-365/developer/publish

Analytics & business value · 8

[60] [Microsoft Copilot agents usage report \(preview\)](#)

Microsoft Learn · Cited on p. 16
learn.microsoft.com/microsoft-365/admin/activity-reports/microsoft-365-copilot-agents-new

[61] [Microsoft Copilot Credits report](#)

Microsoft Learn · Cited on p. 16
learn.microsoft.com/microsoft-365/admin/activity-reports/microsoft-365-copilot-credits

[62] [Connect to the Agent Dashboard \(Viva Insights\)](#)

Microsoft Learn · Cited on p. 16
learn.microsoft.com/viva/insights/org-team-insights/agent-dashboard

[63] [Advanced analysis introduction \(Viva Insights\)](#)

Microsoft Learn · Cited on p. 16
learn.microsoft.com/viva/insights/advanced/analysis-intro

[64] [Copilot Studio agents report \(Viva Insights advanced analysis\)](#)

Microsoft Learn · Cited on p. 16
learn.microsoft.com/viva/insights/advanced/analyst/templates/copilot-studio-agents

[65] [Viva Insights analyst experience \(sign-in required\)](#)

Microsoft Viva Insights · Cited on p. 16
analysis.insights.cloud.microsoft/

[66] [Monitor time and cost savings for agents \(Copilot Studio\)](#)

Microsoft Learn · Cited on p. 16
learn.microsoft.com/microsoft-copilot-studio/analytics-cost-savings

[67] [Communicate the value story to stakeholders \(Copilot Studio\)](#)

Microsoft Learn · Cited on p. 16
learn.microsoft.com/microsoft-copilot-studio/guidance/agent-business-value-tell-value-story

Case studies & announcements · 8

[68] [Implementing Agent 365: How we're governing and managing AI agents at Microsoft](#)

Microsoft Inside Track · Cited on p. 17
www.microsoft.com/insidetrack/blog/implementing-agent-365-how-were-governing-and-managing-ai-agents-at-microsoft/

[69] [KPMG and Microsoft scale trusted, enterprise AI agents globally through deployment of Agent 365 and Copilot](#)

Microsoft Source, June 9, 2026 · Cited on p. 17
news.microsoft.com/source/2026/06/09/kpmg-and-microsoft-scale-trusted-enterprise-ai-agents-globally-through-deployment-of-agent-365-and-copilot/

[70] [Reimagining work, rebuilding value: KPMG leads the AI era with Microsoft-powered agents](#)

Microsoft Customer Stories · Cited on p. 17
www.microsoft.com/en/customers/story/25357-kpmg-azure-ai-foundry

[71] [AGCO scales employee-built AI agents with Microsoft Copilot Studio](#)

Microsoft Customer Stories · Cited on p. 17
www.microsoft.com/en/customers/story/26836-agco-corp-word

[72] [Brown Health scales Microsoft Dragon Copilot and AI agents to ease care delivery](#)

Microsoft Customer Stories · Cited on p. 17
www.microsoft.com/en/customers/story/26765-brown-university-health-microsoft-365-copilot

[73] [HealthEquity scales innovation with Microsoft 365 Copilot and AI agents](#)

Microsoft Customer Stories · Cited on p. 17
www.microsoft.com/en/customers/story/26568-healthequity-microsoft-365-copilot

[74] [EisnerAmper elevates assurance with an Azure AI Foundry-powered agent supporting 18,000 audits](#)

Microsoft Customer Stories · Cited on p. 17
www.microsoft.com/en/customers/story/26018-eisneramper-microsoft-foundry

[75] [Search customer success stories \(Microsoft Copilot - “agent 365”\)](#)

Microsoft Customer Stories · Cited on p. 17
www.microsoft.com/en-us/customers/search?filters=product%3Aai-microsoft-copilot&q=agent+365

About these sources. Links verified September 30, 2026. Microsoft Learn is updated often; confirm preview and Frontier status before implementation. Customer stories illustrate agent rollout patterns; only Microsoft (Inside Track) and KPMG explicitly reference Agent 365. For questions, contact Carlos Suarez (carlossuarez@microsoft.com) or Alejandro Lopez (Alejandro.Lopez@microsoft.com).

Change Log

What changed in each version, where to find it, who updated it and whose feedback shaped it. Newest first.

Version	What changed	Where (current page numbers)	Created or updated by · feedback from
v20260930 September 30, 2026 FastTrack artifact	• Agent 365 Control Plane page (new): admin center areas, agent management rules, lifecycle actions, monitoring signals, admin roles • Purview coverage: what's automatic for agent instances vs. what you must scope (DLP, Insider Risk, Communication Compliance, retention, eDiscovery) • Viva Insights advanced analysis (analysis.insights.cloud.microsoft) for central reporting • Single-tenant Option B: Dev blueprints in the production tenant, isolated with the Environment attribute and a Conditional Access block; suggested Conditional Access policy set • Attribute-based Conditional Access guardrails; “CSA” spelled out as custom security attribute • Cover: date-based version number, owners and contributors; this change log	Agent 365 Control Plane p. 12 Analytics · Usage, Adoption & Impact p. 16 Identity Architecture & Lifecycle p. 5 AuthN / AuthZ & MCP Governance p. 6 Overview p. 1	Updated by: Alejandro Lopez Feedback from: Victor Adeyanju (control plane, Purview, central reporting); customer (single-tenant scenario)
v20260929 September 29, 2026 FastTrack artifact	• Dev · Test · Prod environments: one blueprint per workload per environment, tenant and subscription layout, naming and tagging convention • Environment guidance added for each agent type, the developer workflow, the configuration phases and the operating model (promotion RACI, stale-identity cleanup)	Identity Architecture & Lifecycle p. 5 Agent Types · Identity & Access p. 9 Developer Guide · Pro-Code Agents p. 11 Configuration Blueprint p. 13 Operating Model p. 15	Updated by: Alejandro Lopez Feedback from: Customer (environment strategy question)
v20260928 September 28, 2026 FastTrack artifact	• Published as a Microsoft FastTrack artifact: <i>Microsoft Agent 365 Technical Advisory</i>, built on Carlos's original diagram • Expanded it into an enterprise reference architecture with best practices grounded in Microsoft Learn • Deep dives: identity, authentication and MCP, risk and oversight, security, configuration, connected platforms, operating model • Agent type comparison: Agent Builder, Copilot Studio, Microsoft Foundry and custom pro-code • Technical Advisory Map, analytics, developer guide, case studies, cover page and clickable references • Original Discovery Paths page kept, resized to the document width and its connected-platform list updated	Technical Advisory Map p. 2 Enterprise Reference Architecture p. 3 Agent 365 Discovery Paths p. 4 Agent Types · Risk & Operations p. 10 References p. 18	Updated by: Alejandro Lopez Feedback from: Customer (request for more detail)
Original Original diagram	• Microsoft Agent 365 discovery and inventory diagram: a two-page diagram of how agents are discovered and inventoried in Agent 365 • Its Agent Discovery Paths page is preserved in this document	Agent 365 Discovery Paths p. 4	Created by: Carlos Suarez Feedback from: —

How versions work

- Versions are date based (**vYYYYMMDD**); a second release on the same day adds a suffix, e.g. v20260930.2. The cover shows the current version.
- Every release adds a row here, credits the feedback behind it, and keeps earlier entries unchanged.
- Send feedback to the owners on the cover; contributors are credited on the cover and in this log.