AI AGENT GOVERNANCE

Architecture Design Sessions (ADS) for Agent 365 governance

Interest exists, but no agreed architecture or roadmap

Interest in agents is real, but there is no agreed architecture or roadmap. Competing internal proposals circulate and no target governance model has been settled.

It is still unclear who approves an agent and with what access, and there is no agreed plan for governing agents at scale.

Across the market, the size of that gap is already measurable:

BLIND SPOTS

82%

...discovered unknown agents in the past year

Cloud Security Alliance & Zenity, 2026 (n=445)

ACCESS CREEP

53%

...have had agents exceed their intended permissions

Cloud Security Alliance & Zenity, 2026 (n=445)

POLICY GAP

31%

...have formally adopted an agent policy

Cloud Security Alliance & Zenity, 2026 (n=445)

Before you build, the design has to be agreed.

That is what our E-nvision experience delivers, on predefined high-impact use cases and at the depth your teams need.

Introducing Synergy Advisors Agent 365 E-nvision

E-nvision is the Architecture Design Session (ADS) Synergy Advisors built for this moment: up to three predefined high-impact use cases, based on our experience, discussed in depth with an expert and validated as a high-level design in a functional environment.

It exists because governance decisions taken after deployment are the expensive ones. Built for sponsors, architects and workload owners at Level 300/400, delivered remotely and one to one – the next page details how the sessions run.

- Scenarios and high-level architecture validated in a functional environment, with no changes to your environments.
- Up to three predefined high-impact use cases, discussed in depth with an expert.
- Governance considerations resolved before implementation begins.
- Sponsor, architects and workload owners aligned on one design they agree on.

Agent 365 (A365) – Predefined High-Impact Use Cases

CORE

Observe Agents and Prevent Sprawl
Consolidated view of every agent across platforms, with ownerless and at-risk agents spotted

Onboard Agents with Least-Privilege Access
Approval templates, Entra Access Packages, sponsor validation and Conditional Access

Govern Agent Tools
Govern tools invoked by agents that access different resources or apps, like Outlook and SharePoint, with every action audited

Protect Against Data Oversharing and Risk
Purview DSPM for AI, DLP and Communication Compliance over agent access to sensitive data

Defend Against AI Threats and Vulnerabilities
Prompt injection incidents investigated, compromised agents contained, MCP tools hunted

Audit, Compliance, and Agent Lifecycle
Audit trail exported, non-compliant agents blocked or deleted, records held under retention

Discover and Control Local AI Agents
IDE extensions, CLIs and code editors running as local agents detected on managed devices, unapproved ones blocked with Intune

BETTER TOGETHER

Agent 365 + Defender for Endpoint
Endpoint agents discovered and investigated

Agent 365 + Microsoft Sentinel
Detection, hunting and incident response

Agent 365 + Defender for Cloud
Grounding data and attack path risks

THIRD-PARTY AGENTS

LangChain
Brought under Agent 365 governance and lifecycle

LangGraph
Onboarded with the SDK, same observability and controls

Google ADK
Registered with the SDK, identity, owner, and full observability

Amazon Bedrock AgentCore
Connected platform plus deep integration by our experts

E-nvision (Architecture Design Sessions) Experiences

INFRASTRUCTURE TRACK

Up to **4** SESSIONS — Up to **2h** PER SESSION

Infrastructure Track

- ✓ Up to three use cases (including discovery), discussed with an expert in a functional environment
- ✓ Scenarios, architecture and governance considerations validated interactively
- ✓ Requirements and a validated high-level design of the use case, at Level 300

DEVELOPMENT TRACK

Up to **4** SESSIONS — Up to **2h** PER SESSION

Development Track

- ✓ Up to three use cases (including discovery), focused on third-party agent scenarios with the Agent 365 SDK
- ✓ SDK-based integration, ownership and lifecycle design for external agents
- ✓ Requirements and a validated high-level design of the use case, at Level 400

Scope & Deliverables

Infrastructure Track

IN SCOPE

- ✓ Up to 3 predefined high-impact use cases, based on Synergy Advisors experience
- ✓ Scenario validation in a functional environment
- ✓ Governance and operational alignment
- ✓ High-level architecture reviewed with an expert

OUT OF SCOPE

- ✕ Changes to production or test environments
- ✕ Custom development or deployment tasks

DELIVERABLES

- High-Level Design of the validated use case
- Architecture validation summary
- Governance and next-step recommendations

Development Track

IN SCOPE

- ✓ Up to 3 use cases on third-party agent scenarios
- ✓ Architecture designed with the Agent 365 SDK
- ✓ Integration, ownership and lifecycle design
- ✓ High-level architecture reviewed with an expert

OUT OF SCOPE

- ✕ Changes to production or test environments
- ✕ Technologies outside the workshop objectives

DELIVERABLES

- High-Level Design of the validated use case
- Architecture validation summary
- Governance and next-step recommendations

NEW · THIRD-PARTY AGENTS

Agent 365 beyond Microsoft

AWS, Google Cloud, LangChain, and LangGraph agents, governed in the same Agent 365 control plane as Copilot. Our experts validate and extend what Agent 365 covers out of the box, so every agent meets five conditions.

PLATFORM	DISCOVERY	IDENTITY	SECURITY PROTECTION	DATA PROTECTION	TELEMETRY
Amazon Bedrock AgentCore AWS	Out of the box	Not covered	Partial	Not covered	Partial
Google ADK Gemini Enterprise Agent Platform	Out of the box	Not covered	Partial	Not covered	Partial
LangChain Open source	Not covered	Not covered	Not covered	Not covered	Not covered
LangGraph Open source	Not covered	Not covered	Not covered	Not covered	Not covered
With Synergy Advisors experts Agent 365 SDK · Entra Agent ID Purview · OpenTelemetry	Validated & extended	Validated & extended	Validated & extended	Validated & extended	Validated & extended

Out of the box — Native in Agent 365 (registry sync)
Partial — Read-only signals, no code-level controls
Not covered — Needs Agent 365 SDK integration
Validated & extended — Integrated and tested by our experts

Source: Synergy Advisors, Agent 365 E-xperiences portfolio V.1.7 (slide 47); Microsoft Learn, Connected platforms in Agent 365.

Microsoft delivers, Synergy Advisors extends

Discovery	Every agent registered in Agent 365, with an accountable owner and sponsor.
Identity	Its own Microsoft Entra Agent ID and risk-based Conditional Access.
Security protection	Defender Advanced Hunting and Microsoft Sentinel, one SOC response.
Data protection	Purview DLP checks every prompt and response in real time.
Telemetry	Who uses each agent, how often, and how much, with OpenTelemetry.

On AWS and Google Cloud: from visibility to enforcement

Agents stay where they run. On Google Cloud, ADK callbacks send prompts to Purview.

LEVEL 1

Connected visibility

- ✓ The agent appears in Agent 365
- ✓ Activity synced from Amazon CloudWatch
- ✓ Operations in Defender Advanced Hunting

LEVEL 2

Deep integration

- ✓ Real-time Purview DLP on every prompt
- ✓ Sensitive content blocked
- ✓ DLP evidence in Activity explorer

New · E-mmersion Dev Track

A hands-on lab for developers and platform teams (level 400): onboard LangChain, LangGraph, and Google ADK agents with the Agent 365 SDK, Entra Agent ID, and least-privilege access.

Up to 4 hours (up to 8 in Plus) · In person or remote · English or Spanish · Preconfigured lab, no setup in your tenant

Coming next: more agent platforms and MCP server governance in Azure.

How the sessions run

01 — Preparation
Objectives, stakeholders, and agenda agreed before day one.

02 — Environment Review
Your infrastructure reviewed for gaps and requirements.

03 — Use Case Definition
Predefined scenarios mapped to strategic objectives (up to 3).

04 — Solution Envisioning
High-level architecture validated in a functional environment.

05 — Design Presentation
The proposed design walked through and validated with your teams.

06 — Handover
Requirements, design and deployment recommendations delivered.

A Portfolio of Experiences, One Governance Journey

ALIGN & DESIGN
Create alignment and agree on the target model

- E-nsights
- E-nvision

EXPERIENCE & PROVE
Show it works, then prove it in a dedicated lab

- E-mmersion
- E-xplorer

VALIDATE & DECIDE
Validate functionality before production

- E-valuation

PILOT & SCALE
Move to production and extend governance

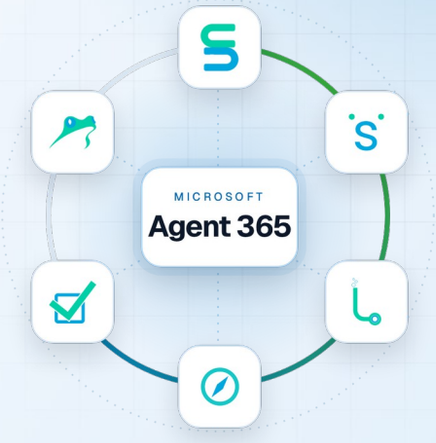
- E-Scuba

Agree the Design Before You Build It

You leave with a target governance model your teams agree on, and a validated high-level design of the use cases you prioritized.

Governance decisions are taken before implementation starts, with effort, sequence and dependencies understood up front.

Let's take the next step together.



For more information, please feel free to contact us at a365@synergyadvisors.biz

Learn more →