SYNERGY ADVISORS

E-xplorer

AI AGENT GOVERNANCE

# Proof of Concept (PoC) in a Box for Agent 365 governance

## A priority scenario chosen, but not yet validated

Teams have chosen a priority scenario but have not yet validated it. In-house agents need structured testing in a lab, and approval stays blocked without hands-on proof of governance.

Teams weighing a pilot need documented evidence they can take into that decision – a record of what was configured, tested and confirmed to work.

**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 commit, the scenario has to be proven.

That is what our E-xplorer experience delivers, on our predefined scenarios and at the depth your teams need.

## Introducing Synergy Advisors Agent 365 E-xplorer

E-xplorer is the Proof of Concept (PoC) in a Box Synergy Advisors built for this moment: a preconfigured, isolated Agent 365 environment, ready from day one, where the predefined scenarios you pick are configured, exercised and functionally validated.

It exists because a demonstration rarely settles a real decision. Governance working on someone else's tenant is not the same as governance holding on yours. Built for technical teams and workload owners at Level 300 (TDM), delivered remotely – the next page details how the sessions run.

- A dedicated lab that is yours alone, explored on your own schedule instead of a single guided session.
- A technical specialist alongside you, resolving questions and validating your design as you go.
- Architecture and functionality confirmed with evidence, before any adoption decision is committed.
- Full freedom to test, break, and reconfigure as often as needed, with no impact on your production tenant.

SYNERGY ADVISORS

E-xplorer

# Agent 365 (A365) – Key Scenarios

## 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-xplorer (Proof of Concept (PoC) in a Box) Experiences

## BASE EXPERIENCE

Up to **2** WEEKS

Up to **2** SCENARIOS

### Base Experience

- ✓ Extended, private lab access in up to 2 weeks, used at your own pace
- ✓ Private sessions with experts to validate design and functionality
- ✓ Up to 2 predefined scenarios validated functionally, not a training course

## EXTENDED SCOPE (PLUS)

Up to **3** WEEKS

Up to **3** SCENARIOS

### Extended Scope (Plus)

- ✓ The same experience, extended up to 3 weeks
- ✓ A third predefined scenario, which may include development
- ✓ Third-party agent and Better Together scenarios included

# Scope & Deliverables

## Base Experience

### IN SCOPE

- ✓ Preconfigured, isolated Agent 365 lab from day one
- ✓ Up to 2 weeks of extended, private lab access
- ✓ Lab setup plus up to 2 predefined scenarios, validated functionally
- ✓ Private sessions with a dedicated technical specialist

### OUT OF SCOPE

- ✕ Any work in your production tenant
- ✕ Client-specific use cases outside the predefined scenario catalog

### DELIVERABLES

- Architecture validation and design recommendations
- Functional validation summary
- Next-step recommendations

## Extended Scope (Plus)

### IN SCOPE

- ✓ Preconfigured, isolated Agent 365 lab from day one
- ✓ Up to 3 weeks of extended, private lab access
- ✓ Up to 3 predefined scenarios, one may include development
- ✓ Third-party agent and Better Together scenarios included

### OUT OF SCOPE

- ✕ Any work in your production tenant
- ✕ Client-specific use cases outside the predefined scenario catalog

### DELIVERABLES

- Architecture validation and design recommendations
- Functional validation summary
- 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**<br>AWS | Out of the box | Not covered | Partial | Not covered | Partial |
| **Google ADK**<br>Gemini Enterprise Agent Platform | Out of the box | Not covered | Partial | Not covered | Partial |
| **LangChain**<br>Open source | Not covered | Not covered | Not covered | Not covered | Not covered |
| **LangGraph**<br>Open source | Not covered | Not covered | Not covered | Not covered | Not covered |
| **With Synergy Advisors experts**<br>Agent 365 SDK · Entra Agent ID<br>Purview · OpenTelemetry | Validated & extended | Validated & extended | Validated & extended | Validated & extended | Validated & extended |

Out of the box — Native in Agent 365 (registry sync)
Not covered — Needs Agent 365 SDK integration
Partial — Read-only signals, no code-level controls
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.

SYNERGY ADVISORS

E-xplorer

## How the sessions run

**01**

### Kickoff

Preconfigured lab ready from day one, predefined scenarios agreed up front.

**02**

### Visibility

Ready-made scenario: discovery, inventory, and shadow AI exposure.

**03**

### Protection

Ready-made scenario: sensitive info types, DLP, and misuse testing.

**04**

### Governance

Ready-made scenario: lifecycle to publish, authorize, and retire.

**05**

### Evidence

Guided hands-on sessions on the predefined scenarios you select.

**06**

### Decision

Validated design and evidence to decide the implementation.

## 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

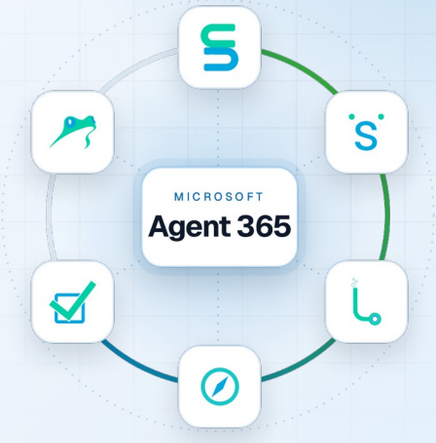
E-Scuba

## Decide With Evidence, Not With Assumptions

You leave with functionality validated on your priority use cases, and the architecture and design confirmed before you commit.

The question is no longer whether this works, but how far to take it.

**Let's take the next step together.**




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

Learn more →