BEST PRACTICES SESSION · 25 MIN

From Chaos to Control: Managing Agents at Scale

The Agent 365 Control Plane

Identity, security, and observability for AI agents - at enterprise scale.

Kranthi Kumar Manchikanti
AI Architect · Microsoft · TheAIRuntime.com

DeveloperWeek New York · AI DevSummit · Expo Stage

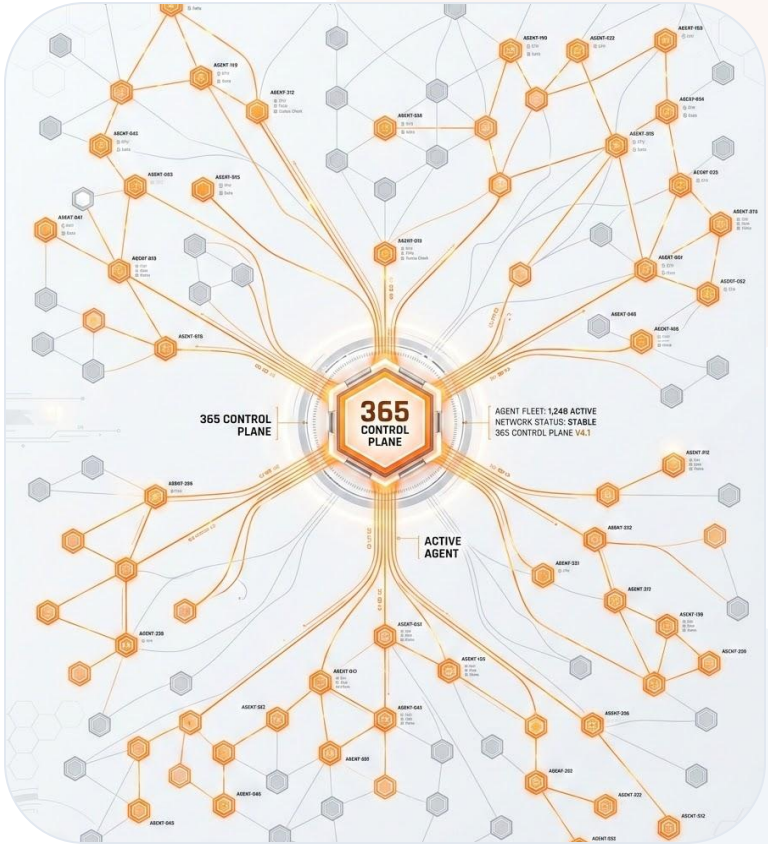

THE SHIFT

Agents are no longer coming. They're already here.

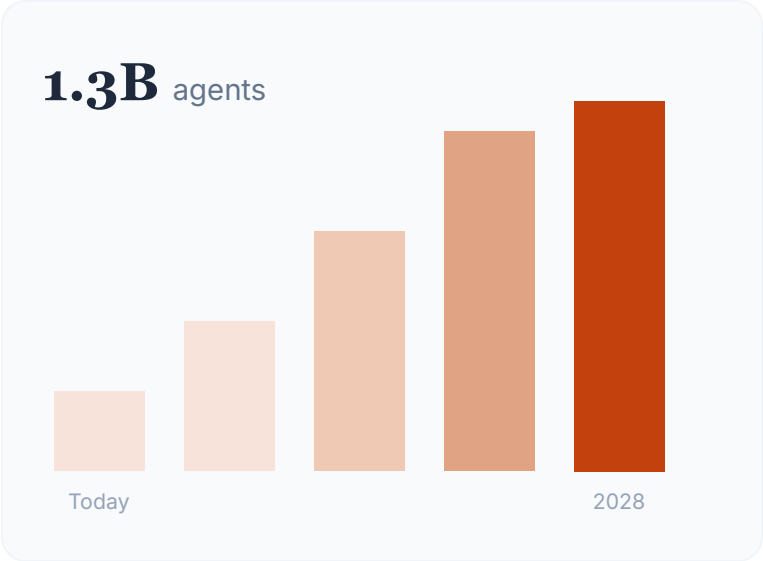

Directional - IDC projects 1.3B AI agents in enterprises by 2028.

They span everything

Copilot, Teams, and M365 - plus local assistants and SaaS agents on emerging platforms.

They act, not just chat

Invoke tools, access data, and talk to other agents - decisions happen in seconds.

They multiply fast

Easier to build and deploy means your attack surface grows with every new agent.

THE PROBLEM

Without a control plane, agents become...

Invisible

No inventory, no owner. Shadow and ownerless agents run outside the view of the teams accountable for risk.

Insecure

Over-privileged access, tool misuse, and prompt injection turn a “helpful” workflow into data oversharing in seconds.

Unmanaged

Built across many tools and frameworks with no consistent policy, lifecycle, or audit trail.

You can't govern what you can't see — and you can't secure what you don't understand.

THE ANSWER

Agent 365 - the control plane for AI agents

Observe, govern, and secure every agent - Microsoft-built, open-source, or third-party - using the admin and security workflows your teams already run.

01

Observe

One inventory of every agent - telemetry, dashboards, and an agent map across the whole fleet.

02

Govern

Identity, least-privilege access, policy templates, and full lifecycle - manage agents like your workforce.

03

Secure

Runtime threat detection, prompt-injection defense, DLP, and compliance built in.

Generally available May 2026 ·

ARCHITECTURE

Built on the systems that already run your enterprise

YOUR AGENTS

Copilot | Teams | Open-source | SaaS agents | Custom

Agent 365 — the control plane: Registry · Access Control · Visualization · Interoperability · Security

Microsoft Entra

Agent ID - identity & access-control foundation.

Microsoft Defender

Threat detection, runtime protection, shadow-AI discovery.

Microsoft Purview

Data security & compliance - DLP, labels, eDiscovery.

Intune + Windows 365

Device discovery and a secured, managed runtime.

No new trust fabric to rebuild — the same identity, security, and compliance stack, now covering agents.

THE CONTROL PLANE

Five capabilities that make enterprise-scale agents possible

The next few minutes: one slide per capability - the field issues each one solves, and what we learned making it work.

CAPABILITY 01 / 5

01

Registry

Stop sprawl at the source: one inventory of every agent.

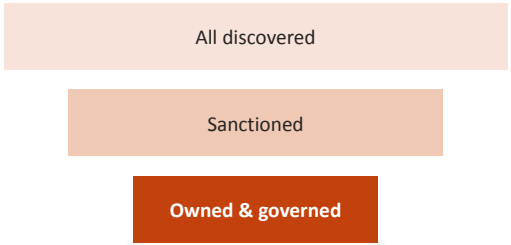

WHAT IT DOES

Single source of truth. Every agent - Entra Agent ID, Teams Store, and discovered shadow agents - in one inventory.
Discovery built in. Defender + Intune find local and cloud agents on the Shadow AI page.
Quarantine on day one. Block unsanctioned agents from connecting to resources.
Multi-cloud reach. Registry sync with AWS Bedrock and Google Cloud (preview).

ISSUES IN THE FIELD

- Shadow agents appear faster than periodic audits can catch them.
- Ownerless agents pile up with no one accountable for risk.
- Multi-cloud agents stay invisible until registry sync is switched on.

WHAT WE LEARNED

- Make discovery continuous - not a one-time audit.
- Require an owner at registration; a registry without owners is just a list.
- Quarantine the unsanctioned first, then onboard the rest under policy.

CAPABILITY 02 / 5

02

Access Control

Every agent gets a unique identity - and only the access it needs.

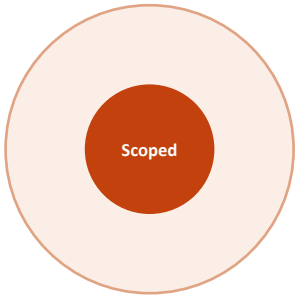

Over-privileged blast radius

WHAT IT DOES

Unique Agent ID. No shared service accounts. Each agent is individually manageable.
Least privilege by default. Scope access to the task; shrink the blast radius from misuse.
Policy templates. Enforce a standard security baseline the moment an agent onboards.
Adaptive Conditional Access. Entra applies risk-based policy in real time and blocks compromised agents.

ISSUES IN THE FIELD

- Shared service accounts erase attribution - you can't tell which agent acted.
- Agents requested broad rights “to be safe,” creating huge blast radii.
- Static grants ignore risk signals and never get walked back.

WHAT WE LEARNED

- One Agent ID per agent, scoped tightly to its task.
- Enforce the baseline with policy templates at onboarding, not later.
- Let Conditional Access revoke access automatically on risk.

CAPABILITY 03 / 5

03

Visualization

Move from monitoring to actionable insight across the fleet.

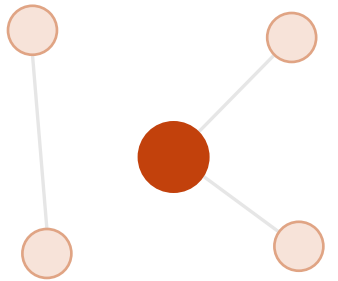

Agent Map — agents · users · resources

WHAT IT DOES

Agent Map. A complete map of connections among agents, users, and resources.
Telemetry & alerts. Unified dashboards with anomaly detection so blind spots don't become incidents.
Role-based reporting. IT, security, and business leaders each see the metrics that matter.
ROI & compliance. Per-agent performance plus logging and eDiscovery to stay audit-ready.

ISSUES IN THE FIELD

- Raw telemetry is just noise without a baseline of “normal.”
- Dashboards exist but no team owns or watches them.
- Each audience wants a different cut of the same data.

WHAT WE LEARNED

- Define alerts and ROI metrics before go-live, not after.
- Use the Agent Map to trace every agent-to-data path.
- Give IT, security, and business their own role-based views.

CAPABILITY 04 / 5

04

Interoperability

Agents get the same context as the people they work alongside.

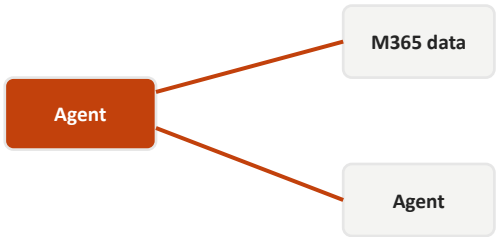

Task-scoped, not standing access

WHAT IT DOES

Native M365 access. Word, Excel, SharePoint, Dynamics 365 - the same data your users touch.
Framework-agnostic. Microsoft, open-source, and third-party agents are first-class citizens.
Agent-to-agent. Secure interoperability lets agents collaborate without bespoke plumbing.
Three operating modes. Delegated and behind-the-scenes agents are GA; team workflows in preview.

ISSUES IN THE FIELD

- Broad standing permissions get granted just “to make it work.”
- Agent-to-agent calls happen without any scoping.
- Framework sprawl makes consistent policy hard to enforce.

WHAT WE LEARNED

- Grant task-scoped access through the Agent ID - never standing broad rights.
- Treat interop as a privilege to scope, not a door to leave open.
- Make OSS and third-party agents first-class, but governed.

CAPABILITY 05 / 5

05

Security

Defense in depth — from runtime behavior to the data agents touch.

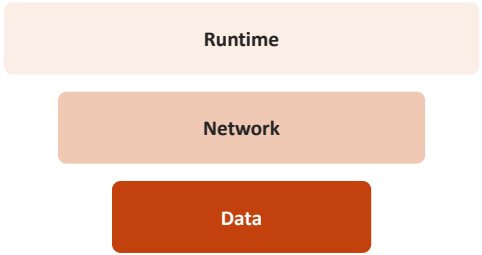

Defense in depth

WHAT IT DOES

Runtime defense. Defender XDR detects threats and protects agents as they execute.
Prompt-injection blocking. Network-level protection stops injection before it reaches the agent.
Data security. Purview DSPM for AI, DLP, and sensitivity labels guard against oversharing.
Insider risk & audit. Detect, retain, and investigate risky agent interactions.

ISSUES IN THE FIELD

- Prompt injection and tool misuse turn helpful flows into data leaks.
- A single guardrail fails silently and you never know.
- Write-access to core systems raises the stakes of every mistake.

WHAT WE LEARNED

- Layer runtime + network + data controls - one guardrail isn't enough.
- Assume the agent is a target and plan blast-radius containment.
- Use Purview DSPM/DLP to catch oversharing before it ships.

BEST PRACTICES

Patterns for safely scaling agent-based solutions

01 Manage agents like a workforce

Onboard, assign an owner, and offboard. No ownerless agents.

02 Identity-first, least privilege

A unique Agent ID per agent, scoped to its task. Templates enforce the baseline.

03 Discover before you govern

Switch on shadow-AI discovery early; quarantine, then bring the rest under policy.

04 Instrument for observability

Define alerts and ROI metrics at launch. Watch agent-to-resource paths.

05 Layer your defenses

Runtime + network + data controls together. Assume compromise.

06 Reuse the trust fabric

Extend Entra, Defender, and Purview — don't rebuild governance.

ADOPTION ROADMAP

From where you are to enterprise AI at scale

01

Crawl

Discover & inventory

- Enable registry + shadow-AI discovery
- Baseline every agent and its owner
- Quarantine the unsanctioned

02

Walk

Govern

- Apply policy templates & least privilege
- Adaptive Conditional Access via Entra
- Stand up dashboards & alerts

03

Run

Secure & scale

- Runtime + prompt-injection defense
- Purview DLP, labels & eDiscovery
- Measure ROI; sync multi-cloud agents

Microsoft Agent 365
CASE STUDY

BIOTECHNOLOGY & R&D

Amgen

From concept to a governed R&D agent in just six weeks.

THE CHALLENGE
Speed up drug research without adding new governance or security risk.

THE AGENT
An R&D agent that mines research data and drafts scientific insights.

GOVERNED BY AGENT 365
A unique Entra Agent ID, least-privilege access, and full auditability.

Source: Microsoft customer stories · Microsoft Ignite 2025

OUTCOME

6

WEEKS

from idea to a production-ready, governed agent.

- Built with Microsoft Copilot Studio
- Observed & secured end-to-end in Agent 365

PROJECTED IMPACT

Millions

in targeted cost savings.

- Across global productivity and supply chain
- Powered by Microsoft 365 Copilot and agents
- Every agent carries a governed identity

Microsoft Agent 365
CASE STUDY

MATERIALS SCIENCE

Dow

Reimagining productivity and supply chain with a fleet of agents.

THE CHALLENGE
Scale agents across global operations without losing oversight.

THE AGENTS
Productivity and supply-chain agents working alongside employees.

GOVERNED BY AGENT 365
One control plane to observe, secure, and measure the whole fleet.

Source: Microsoft customer stories · Microsoft Ignite 2025

TAKEAWAYS

If you remember three things

01

Visibility is the prerequisite.

You can't govern or secure what you can't see. Start with discovery and a single registry.

02

Identity is the control point.

A unique Agent ID with least-privilege access turns ungoverned agents into managed ones.

03

Extend, don't rebuild.

Agent 365 brings agents into Entra, Defender, and Purview — the trust fabric you already run.

Thank you.

Let's govern the agent fleet - before it governs us.

Kranthi Kumar Manchikanti
AI Architect · Microsoft · TheAIRuntime.com

Go deeper aka.ms/Agent365 · Agent Governance Whitepaper

Questions? Find me at the Expo Stage.

Scan to revisit on your phone.

events.theairuntime.com/agent-365