

AI Agent Governance Starter Kit: Rovo and Copilot Studio

Where to build your first enterprise agent, and how to govern both platforms with one model

Most enterprises now have two places to build AI agents: Rovo Studio inside Atlassian and Copilot Studio inside Microsoft. Both are already being used by teams whether or not IT has a policy. This sheet helps you decide where the first governed agent should live, and how to govern both with one model.

What each platform is good at

	Rovo Studio (Atlassian)	Copilot Studio (Microsoft)
Native data	Jira, Confluence, JSM, Bitbucket, and the Teamwork Graph, plus connectors	Microsoft 365 content, Dataverse, Dynamics 365, and Microsoft Graph, plus connectors
Best first use cases	Backlog grooming, ticket triage and Level 1 support (Rovo Service), release notes, PI planning prep, knowledge answers from Confluence	Employee self-service on HR and IT policy, document drafting, meeting follow-ups, process automation with Power Automate
Agent controls	Roles, approvals, versioning, audit logging, org-wide agent inventory, separate permissions for using AI vs. building agents	Lifecycle and evaluation controls, multi-agent orchestration, MCP connectors, Agent 365 inventory and cost visibility
Who builds	Jira and Confluence admins, delivery leads	Power Platform makers, IT, business analysts

Pick the first agent by the workflow, not the vendor

If the workflow lives in tickets and pages, build in Rovo. If it lives in documents, email, and Teams, build in Copilot Studio. The first governed agent should have a clear owner, a bounded data set, a human approval step, and a measurable outcome. A JSM Level 1 deflection agent and an HR policy assistant are the two most common starting points.

Govern both with one model

- ☐ **One inventory.** Every agent, on either platform, is registered with an owner, a purpose, a data scope, and a risk tier. The governance record sits above the platform inventories.
- ☐ **One set of decision rights.** Who may build, who approves data access, who supervises autonomous actions, and who can retire an agent. Document it as a RACI and enforce it in platform permissions.
- ☐ **One control map.** NIST AI RMF functions (Govern, Map, Measure, Manage) mapped to platform controls and to the regulations you already answer to: HIPAA, model-risk guidance, state AI laws. ISO/IEC 42001 where certification is the goal.
- ☐ **One measurement.** Adoption, quality (sampled accuracy and escalation rate), and control metrics (agents without owners, agents touching sensitive data) reported monthly.
- ☐ **One change program.** Roles change when agents take work. Training, communication, and support are the same program regardless of where the agent runs.

AI Agent Governance Starter Kit: Rovo and Copilot Studio

Where to build your first enterprise agent, and how to govern both platforms with one model

Where the platforms meet

Rovo can read Microsoft 365 content through connectors, and Copilot Studio can call Jira through connectors and MCP. The integration question is less about data access and more about which system of record owns the action. An incident is owned in JSM; an approval memo is owned in SharePoint. Agents can cross the boundary, but the record should not.

Talk to a practitioner

VTPMO builds one AI governance model across Rovo and Copilot for regulated organizations: agent inventories, decision rights, NIST AI RMF and ISO 42001 alignment, and the training that makes adoption stick.

Book a 30-minute discovery call: outlook.office.com/book/vtpmo@vtpmo.com · info@vtpmo.com