



Copilot4DevOps

V9 (Standalone Version)

RELEASE NOTES





TABLE OF CONTENTS

General	3
New Feature	3
Bug Fixes	3
Admin	5
Tool Enhancement	5
Bug Fixes	5
Chat	6
New Feature	6
Tool Enhancement	7
Bug Fixes	9
Elicit	10
Tool Enhancement	10
Bug Fixes	10
AI Edit	11
Tool Enhancement	11
Bug Fixes	12
Analyze	13
Bug Fixes	13
SOP/Document Generator	14
Bug Fixes	14
Convert	15
Bug Fixes	15
Transform	16
Bug Fixes	16
Dynamic Prompt	17

Bug Fixes 17

AGENTS4DEVOPS18

New Feature 18

GENERAL

NEW FEATURE

1. Migration from Prompt-Based Execution to Agent-Based Architecture

As part of the ongoing platform architecture improvements, the execution model for all the Copilot4DevOps functionalities has been migrated from a prompt-driven approach to an agent-based architecture.

Previous Behavior

Previously, functionality was managed through system prompts, where dedicated prompts were passed directly to control:

- Response structure,
- Execution behavior,
- Function handling,
- Conversation orchestration

What's new

Copilot4DevOps now runs on a new agent-based architecture that replaces our previous prompt-based system.

In practice, this means:

- **More reliable responses** across all Copilot4DevOps functions, with centralized execution and consistent behavior
- **Improved chat experience** chat execution is now managed through agents (by default set the worker agent with related skills to generate the response), with greater flexibility for complex conversations
- **A more scalable platform** that allows us to roll out new capabilities faster

Note: Nothing changes in how you use Copilot4DevOps. All existing features, workflows, and interactions remain the same you'll simply benefit from improved reliability and performance behind the scenes.

2. Multimodal AI Support for Images in Documents and Work Items

Copilot4DevOps now supports end-to-end image processing and rendering across uploaded documents, BYOD content, work item attachments, and Azure DevOps work item fields. Images embedded within descriptions, acceptance criteria, repro steps, comments, and other HTML fields can now be detected, processed, preserved, and displayed directly within the Copilot4DevOps UI as part of AI operations—enabling richer multimodal (text + image) experiences.

Extracted images are maintained with their contextual placement, surrounding content, and captions where available, while the Copilot4DevOps UI renders embedded images inline for a more accurate and visually consistent experience. These enhancements enable LLMs to process both textual and visual information together, improving AI understanding and supporting documentation-heavy enterprise workflows more effectively.

BUG FIXES

1. Fixed an issue where SOP/Elicit were not processing the entire uploaded document when generating the output.

2. Fixed an issue where some Wiki pages and sub-pages were missing from the BYOD tab.
3. Fixed an issue where a continuous loader was displayed while rewriting the prompt using “Rewrite Prompt” option in a Text Block
4. Fixed an issue where Chat and other modules failed to work when Copilot4DevOps was accessed from the hub on an ADO Server running a local extension, causing authentication errors in Chat and preventing work items from loading in other modules.


ADMIN PANEL

TOOL ENHANCEMENT

1. Admin Control for Automatic Comments in Work Items

We've introduced a new **admin-level configuration** that allows organizations to control whether system-generated comments are automatically added to work items during **publish and update actions**.

- A new checkbox has been added under **Admin Panel → General Tab: “Allow Copilot4DevOps to add comments into work items during publish and update”**.
- The existing tag section label has been updated to reflect control over both **tags and comment**
- If **enabled** Copilot4DevOps continues to automatically add comments during **Work item creation (publish)** and **Work item updates**
- If **disabled**, Copilot4DevOps will not add any automatic comments across supported features.



BUG FIXES

1. Fixed an issue

CHAT

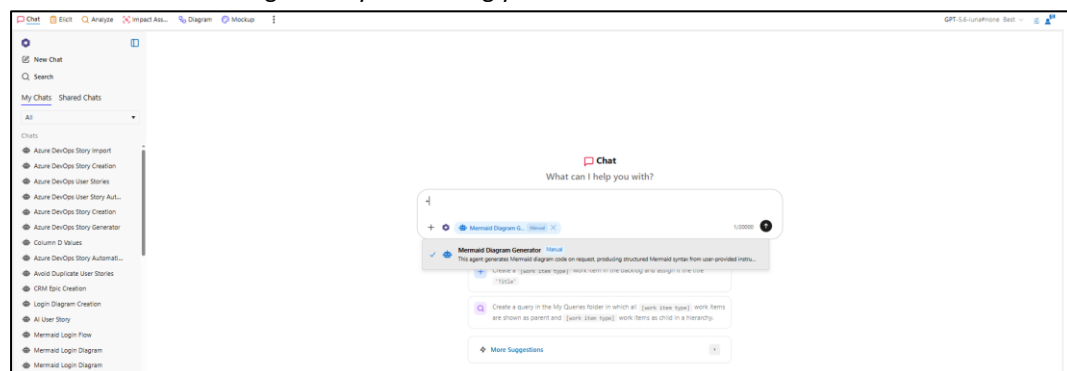
NEW FEATURE

1. Agents with Chat Prompt

We've introduced Agents in Chat, allowing you to directly select specialized AI agents from the prompt to tailor responses based on your task. Agents act as **specialized AI assistants**, each designed for a specific purpose such as generating queries, analyzing work items, etc. By selecting an agent, you guide the AI to respond with the right context, structure, and intent, helping you get more accurate and relevant results.

What you can do

- Select an agent using the **“+” button** or by typing **“+”** in the prompt
- The selected agent appears as a **chip** and stays active until removed or replaced
- Switch or remove agents anytime during your conversation



a. Applicable to Both Single-Agent and Multi-Agent Modes

i. Worker Agent:

Worker Agent is the default fallback execution agent used when:

- No agent is explicitly selected, or
- Dynamic Agent Selection cannot find a relevant custom agent, or
- The Dynamic Agent Selection toggle is OFF (meaning the system will not attempt to auto-pick custom agents).

In other words, Worker Agent ensures the system can still execute end-to-end tasks reliably even when agent selection is not configured, disabled, or yields no match.

ii. Dynamic Agent Selection:

When the Dynamic Agent Selection toggle is ON, the system:

- Analyzes the user request intent and attempts to auto select the most relevant agent(s) from the available set (custom agents if present; otherwise fallback behavior applies).
- Uses those selected agents for execution that match their skills/tools.

If Dynamic Agent Selection is ON but no suitable custom/specialized agent is found, the system falls back to the Worker Agent to avoid blocking execution.

b. Agent behavior in Chat:

The experience for chat using agent is available directly within Copilot4DevOps, providing focused, context-aware assistance for single agent.

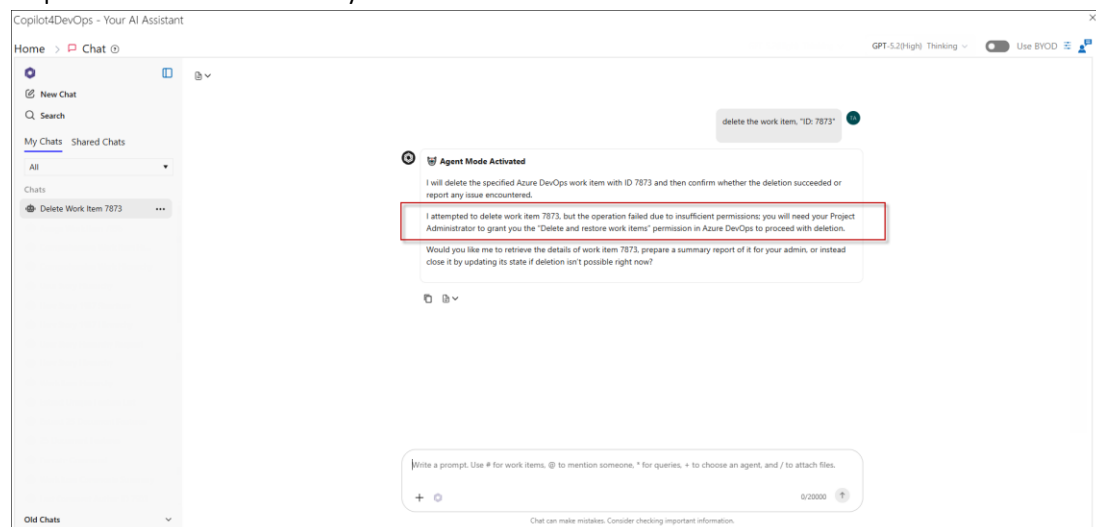
- Allows users to ask natural-language questions about DevOps work, such as interpreting work items, pipelines, or artifacts, and receiving clear, actionable responses in one place.
 - Provides task-oriented guidance without requiring users to leave their existing Azure DevOps context.
 - Maintains short-term conversational context so follow-up questions are answered consistently and efficiently.
 - Designed for simplicity and speed, making it ideal for common day-to-day productivity tasks handled by a single agent.
- i. **Agent referencing:**
- In Chat, exactly one agent is selected and referenced for the full request.
 - If the user explicitly selects an agent, that agent is used for the entire execution. If not, the system uses the default Worker Agent depending on configuration and availability.

TOOL ENHANCEMENT

1. Improved Permission Error Messages

We've enhanced permission error messages in Copilot4DevOps to make them clearer and more helpful. Previously, users received **generic error messages** when an action was restricted, making it difficult to understand the cause. With this update, error messages now clearly explain **what action failed, why it was blocked, and what needs to be done to resolve it**.


This improvement helps users **quickly identify permission** issues, reduce confusion, and take the right steps to resolve them efficiently.



2. Improved Chat Prompt Menu


The chat prompt interface has been enhanced to provide a more organized and flexible user experience.

- The existing **“+” menu** has been expanded:
 - Previously used only for file uploads
 - Now also includes:
 - **Insert Text Block**
 - **All Agents** selection



3. Add a comment into the work item

Whenever a work item is updated through Chat, the system now automatically adds a comment Updated from Copilot4DevOps Chat to the work item to clearly indicate that the update was performed via Chat for better traceability and audit history.



4. Improved Work Item Retrieval

This release improves how work items are retrieved by removing the previous fixed pagination behavior. The system now dynamically returns work items based on the request and available results.

- If the result set is small, it may return all work items from the list or query.
- Users can also explicitly request a specific number of work items (for example, 100, 200, or other), and the system will return that exact number whenever sufficient results are available.

Note: For very large result sets (for example, queries returning thousands of work items), the number of work items returned may vary when no specific count is requested. To retrieve an exact number of work

items, include the desired count in your prompt but do not mention a higher count. It's possible that the LLM may not send a higher count at once.

BUG FIXES

1. Fixed an issue where the diagram was not generating in chat feature when using LLM service **"AzureAI"**
2. Fixed an issue where the image was not displayed in the work item when adding through chat.
3. Fixed an issue where test cases generated from a work item via Chat were being linked as **Child** or **Related** instead of **Tested By**
4. Fixed an issue where downloaded Excel files were being corrupted.
5. Fixed an issue where the **Save As** icon was not visible in the Chat right panel.
6. Fixed an issue where a comment was being added to the Discussion tab when fetching work items in Chat, even when the work item had been created manually.

ELICIT

TOOL ENHANCEMENT

1. Enhanced Elicit Workflow with Guided Next Work Item Suggestions

The Elicit experience has been enhanced to help users stay in flow after creating a work item. Once a work item is published, Copilot4DevOps intelligently suggests the **next relevant work item types** based on hierarchy and project configuration, allowing users to seamlessly continue building their backlog without interruption.

These guided suggestions appear beside the **Generate button** and adapt to the current context, enabling quick continuation of work item creation. Copilot4DevOps also remembers user selections, ensuring consistency and reducing repeated setup, making the overall elicitation process faster, more intuitive, and aligned with real-world workflow progression.

2. Enhanced Work Item Creation with Intelligent Area Path Handling

Copilot4DevOps now automatically inherits the Area Path from the parent work item when creating child work items, ensuring better alignment with Azure DevOps project structure and organizational standards. Before creating the work item, the system validates whether the user has the required permissions for the inherited Area Path.

If the user has access, the work item is created successfully with the inherited Area Path. If access is restricted, Copilot4DevOps provides guided handling based on the access context. In standalone Copilot4DevOps, the Work Item window opens, allowing users to modify the Area Path before saving. When Copilot4DevOps is accessed from ModernRequirements4DevOps, an Area Path dropdown is displayed on the right panel, enabling users to select an accessible Area Path and continue work item creation without interruption.

BUG FIXES

1. Fixed an issue where title generation in Elicit stalled at 90% and failed to complete when only a detailed description was provided.
2. Fixed an issue where, after a session expired, files re-uploaded in Setup Elicitation continued to be reported as unavailable, even following a browser refresh.

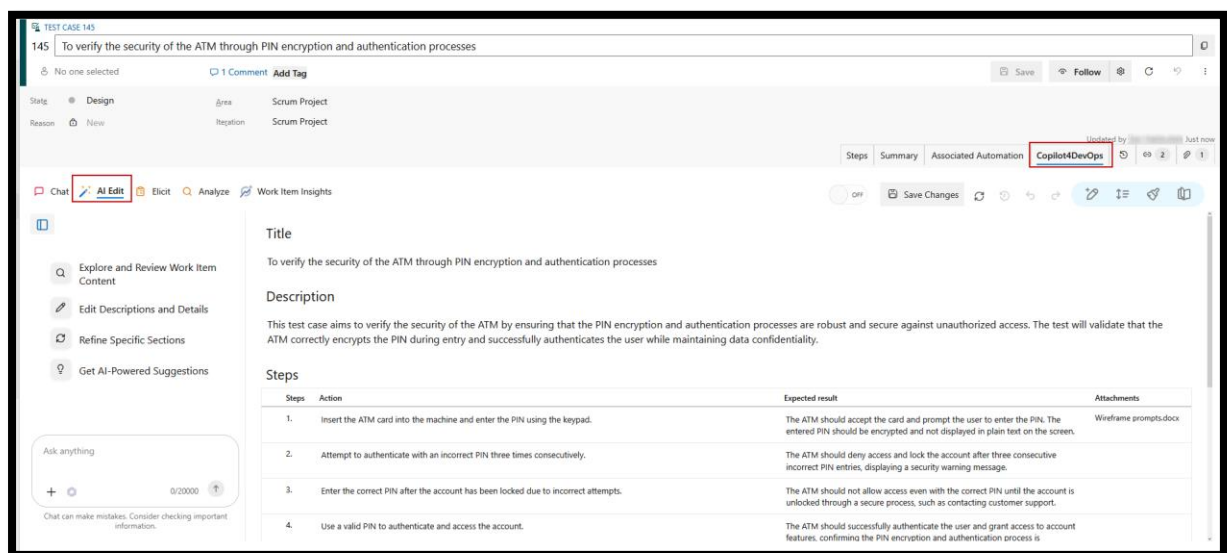
AI EDIT

TOOL ENHANCEMENT

1. AI Edit Support Added to the Copilot4DevOps Work Item Tab

The Copilot4DevOps Work Item tab has been enhanced with integrated AI Edit support, enabling users to refine and update work item content directly within the Azure DevOps work item experience. Alongside Copilot4DevOps Chat, Elicit, Analyze, and Work Item Insight, users can now leverage AI-powered editing capabilities without leaving the work item context—reducing context switching and improving productivity.

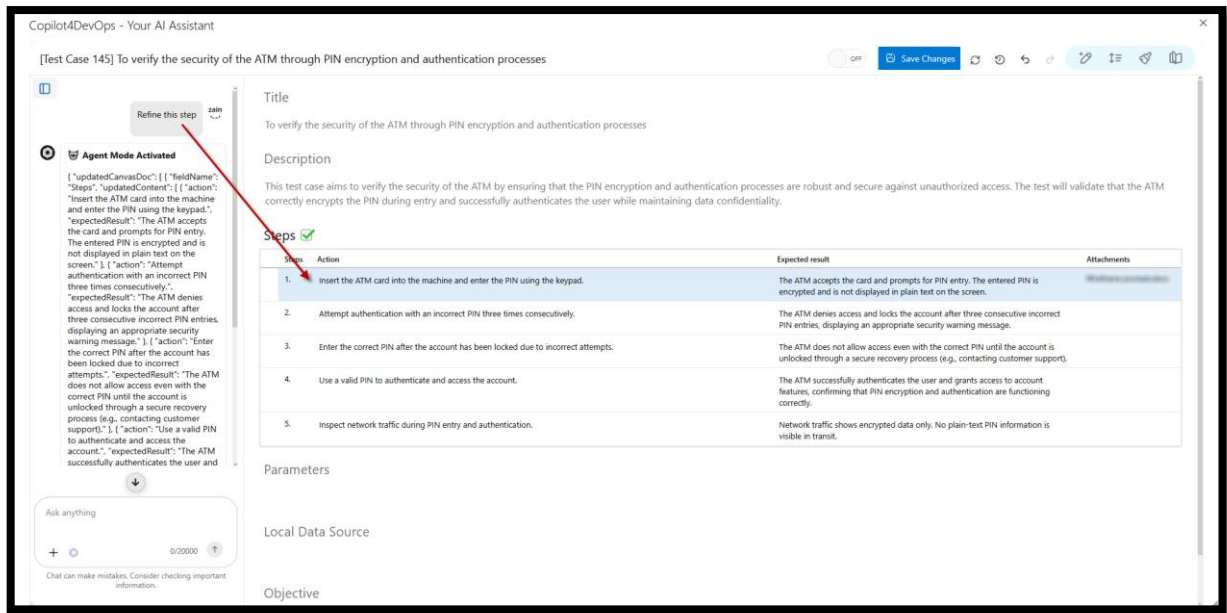
AI Edit sessions are designed for lightweight, in-context interactions, and chat history is not retained when users navigate to another work item tab. Feature availability continues to respect license-based access, ensuring a consistent and streamlined experience across supported environments.



2. Inline and Chat-Based Editing Support for Test Steps in AI Edit

AI Edit now supports editing of Test Steps tables directly within the Copilot4DevOps experience. Users can refine and update test steps through both inline editing and AI-assisted chat interactions, while preserving the structured table format for actions, expected results, attachments, and related test content.

This enhancement improves usability for test management workflows by enabling faster updates, easier refinement of test scenarios, and a more seamless AI-assisted editing experience within Azure DevOps work items.



BUG FIXES

1. Fixed an issue where custom fields added to work items through a custom process template were not displayed in AI Edit.

ANALYZE

BUG FIXES

1. Fixed an issue where the response UI is not proper. Some elements that were previously misaligned are now properly aligned and positioned.
2. Fixed an issue where updating a work item through the Chat Analyze feature was adding multiple comments instead of a single comment.

SOP/DOCUMENT GENERATOR

BUG FIXES

1. Fixed an issue where generated documents were not consistently following the formatting defined in the Word template, including styles, headings, layout, and overall structure.
2. Fixed an issue where long WIQL queries were failing in the SOP/Document Generator in production.

CONVERT

BUG FIXES

1. Fixed an issue where Gherkin output generated by the Convert function was not properly formatted.

TRANSFORM

BUG FIXES

1. Fixed an issue where the Copy option remained disabled after a response was successfully generated in the Transform module.

DYNAMIC PROMPT

BUG FIXES

1. Fixed an issue where opening a saved prompt from the Dynamic Prompt caused a continuous loading state, and the prompt failed to save.
2. Fixed an issue where running a Dynamic Prompt across large work item queries caused excessive memory consumption and led to a crash.

AGENTS4DEVOPS

NEW FEATURE

Agents4DevOps delivers complete, end-to-end experience for configuring AI agents, managing their permissions, running them across projects, and monitoring their execution history. It also introduces a reusable skills library and rich administration controls so teams can safely adopt automation while retaining clear governance and visibility.

Agents4DevOps Capabilities:

1. Jobs Tab – Centralized Agent Execution History

The Jobs tab now serves as the master execution history for all agents within a project, giving users a clear, filterable view of how agents are running overtime. This makes it easier to monitor automation, identify failures, and understand where user input is needed.

- See every agent run as a job in a dedicated grid, with columns for Job ID, Agent name, trigger event, work item type, repository, status (In Progress, Success, Failed, Terminated), number of open tasks, and timestamp.
- Filter jobs using a flexible filter bar with options for agent name, trigger event, work item type, repository, and status, and a dedicated “Open Tasks” filter that shows only jobs waiting for user input.
- Refresh the grid to pick up new jobs that start while the user is on the page and use row-level checkboxes plus a bulk Delete action (with confirmation) to remove historical jobs that are no longer needed.
- The Usage view opens a dedicated Credits Usage dashboard that gives users a complete breakdown of how credits are consumed across all agents in a project over a selected time period, making it easy to track spend, pinpoint high-consumption agents, and understand execution patterns at both a summary and transaction level.
- The Credits Status option provides an at-a-glance summary of the organization's credit balance directly from the Jobs tab, so users can quickly check where they stand without opening the full Usage dashboard.
- Access a context menu for each job with options to view details in Job Explorer, terminate an in-progress job, or delete it, ensuring users can intervene when executions need to be stopped or cleaned up.

Note: User can also access the Jobs tab by clicking on an agent name from the Agents tab. All details will remain the same as master job tab except Usage and credit status options; however, this view will display only the jobs filtered for the selected agent.

a. Job Explorer – Deep Dive into Individual Job Runs

For each job, the Job Explorer provides a rich, three-tab experience that surfaces execution details, agent responses, and where applicable interactive tasks that require user decisions. This helps users understand what happened, why, and what to do next.

- Use the Run Details tab to see the job's status, core information (Job ID, and for event-based runs, the associated work item type and trigger event), execution model, skills, execution timing (start, end, duration), and logs.
- Review logs that show high-level activity for manual and scheduled runs, and deeper, step-by-step processing for event-based runs, including configuration resolution, code execution, and input/output details that support troubleshooting and validation.
- Use the View Agent Response tab to see the agent's own narrative output, including its plan, findings, and any limitations or errors, presented in a human-readable format that explains what the agent attempted to do.
- For assisted autonomous agents that require approvals, use the Tasks tab to work through pending questions: see a highlighted message explaining that the agent is waiting for input, review each task in a card layout with status, question, and answer field, and progress through tasks using Next/Proceed buttons until execution can continue.
- User can get help from an integrated Chat Assistant within Job Explorer to ask questions about the agent's execution details, get help to interpret tasks and guide users by providing appropriate responses to the questions.

2. Agents Workspace – Viewing and Managing Agents

Users can now open the Agents tab in Agents4DevOps to see all configured agents in a single, filterable grid, making it easy to understand what automation is available in a project or across the organization. The grid clearly distinguishes between system and custom agents with collapsable sections, shows how each agent is configured, and allows users to quickly inspect or act on any agent.

- View all agents in a grid with key details such as name, configuration including the agent mode and it's type, project name, trigger event and enabled status.
- Use rich filtering to narrow agents by name, mode, type, projects, trigger event and enabled status, with multi-select dropdowns.
- Interact with a contextual menu per agent to view details, clone an existing configuration, edit agent, execute assisted autonomous agents, or permanently delete agents with confirmation safeguards.

a. Creating Execution Agents from Instructions, Code, or Templates

Execution agents can now be created through multiple flexible paths, allowing teams to either describe the behavior they want, upload existing code or configuration, or reuse prebuilt templates. This reduces setup time and ensures that automation can be tailored to each project's needs.

- Use the “Execution Agent” mode from the New Agent flow and select between three creation methods cards: writing instructions, uploading configuration/code, or starting from a library template.
- In the “Write Instructions” method, specify a code generation model, select one or more projects, define work-item-based triggers and work item types, and provide natural-language instructions. The system then generates agent code, name, description, and tags that can be reviewed and refined before creation.
- In the “Upload config or code” mode, upload existing source code or configuration files (such as MatCal or suspect link configurations) and let the system drive trigger events, work item types, and code details automatically, with the ability to adjust them prior to creating the agent.
- In the “Library” mode, browse a curated library of execution agent templates, open any template to review its description and tags, and then “Use” it to create a new agent that can be customized before being added to the grid.
- Access version management for skills, allowing users to see previous versions, preview older definitions, and restore them when needed, with clear version labels, timestamps, and creator or updater information.
- User can get help from an integrated Chat Assistant available on each configuration tab to ask questions about the current agent, and to update code according to requirements.

b. Creating Assisted Autonomous Agents

The product now supports creating assisted autonomous agents through a guided configuration experience that balances automation with human oversight. Users can define how and when these agents run, whether they require approvals, and how they should orchestrate work.

- Start from the Agents tab using the “New Agent” button, choose “Assisted Autonomous Agent” as the mode.
- User can also choose from multiple creation modes: ‘Create with AI’ (describe the goal and let the system draft detailed agent instructions and regenerating on different models), ‘Create manually’ (you define name, description, and instructions), or starting from ‘Library’ templates that can be tailored to your needs.
- Configure the scope of each agent by selecting one or more projects (including an “ALL” option), choosing an execution model from available LLMs, and optionally require approval so that the agent pauses and creates tasks when user input or clarification is needed.
- Select from Manual, Schedule, or Event-based agent types, with tailored options such as once, daily, weekly, monthly, or yearly schedules, or a rich set of DevOps event triggers spanning builds, code, pipelines, and work items.

- For event-based agents, define exactly which events and work item types or repositories should trigger execution, with support for multi-select options so agents can react to the right changes across code, pipelines, and work items.
- Use structured configuration tabs to manage behavior:
 - Trigger tab for all timing and event settings (including Required Approval, projects, and execution model)
 - Details tab for name, description, and instructions
 - Tools tab to search, filter, and add skills or use AI-driven skill suggestions, as well as manage tool calls by enable/disable them from provided toggles for the specified agent
 - Notifications tab to turn notifications on/off, configure subject/body, include logs and agent responses, and select recipients.
- Access version management for agents, allowing users to see previous versions, preview older definitions, and restore them when needed, with clear version labels, timestamps, and creator or updater information.
- User can get help from an integrated Chat Assistant available on each configuration tab to ask questions about the current agent, and receive guidance when choosing tools or adjusting settings.

3. Skills Tab – Viewing and Managing Skills

A dedicated Skills tab now provides a centralized view of all skills defined in the organization, helping teams understand and govern reusable capabilities that agents can use. The grid clearly distinguishes between system and custom skills and shows where each skill is used.

- View all skills in a grid distributed in 2 collatable categories System and Custom, in which skills exist with columns for name, linked agents, and enabled status, so you can quickly see which skills are active and where they are applied.
- Filter skills by name, category, and linked agents using a rich filter bar with multi-select dropdowns and quick “clear” actions, making it easy to find the right skills in large environments.
- Use the “New Skill” button to create new custom skills and manage existing ones through a context menu that supports cloning, editing, and deleting, with system skills restricted to Customize and Clone to protect core functionality.
- Enable or disable skills via a toggle; when a skill is disabled, it remains visible but is not used during agent execution, even if it is configured on an agent, ensuring safe deactivation without losing configuration history.
- Access version management for skills, allowing users to see previous versions, preview older definitions, and restore them when needed, with clear version labels, timestamps, and creator or updater information.

- User can get help from an integrated Chat Assistant available on each tab to ask questions about the skill, perform update as well as can get information about the artifacts attached as knowledge for the skill.

4. Library – Reusing Pre-Created Agents and Skills

The library tab introduces a curated catalogue of pre-created assisted autonomous agents, execution agents, and skills that can be discovered, filtered, and reused. This helps teams start from proven patterns instead of designing every agent or skill from scratch.

- Browse three library sub-tabs—Assisted Autonomous Agents, Execution Agents, and Skills—each with its own search and collapsible categories distribution.
- Search by keyword across names, and descriptions is available and user can see category labels with counts to understand how many templates are available in each area.
- Open any library item in assisted autonomous agent tab, to see a detailed popup with fields such as name, library category, description, and tools (for assisted autonomous agents), along with a prominent “Use” action.
- Click “Use” to open the selected agent or skill in an editable form, with a clear notice that it is based on a predefined process template and should be validated for the current project, then create it so it appears in the main Agents or Skills grid.

5. Configuration Agents4DevOps

a. Agent Connection Settings

Administrators can now configure how Agents4DevOps connects to backend services, including personal access tokens and base URLs. This ensures that agents operate securely and consistently across environments.

- Use the “Configure Agents4DevOps” section under the General tab in the Admin Panel to define agent connection settings, including one or more personal access tokens (PATs) and the base URL for the Copilot4DevOps service.
- Enter a primary PAT, validate it with a dedicated button, and see the associated PAT user to confirm that the correct identity is being used.
- Add up to ten PATs, including the “Additional PAT” field, so that if the primary PAT becomes invalid, the system can fall back to additional tokens without disrupting agent operations.
- Configure or override the default base URL (prepopulated with the hosted service address) when using a self-hosted or alternative deployment of Copilot4DevOps, ensuring agents connect to the correct backend.

b. Outreach Email Configuration

Administrators can now configure how Agents4DevOps sends outbound email notifications based on agent instructions, using either an SMTP server or the SendGrid Web API. This provides centralized, organization-wide control over agent email and ensures that outreach is sent from an approved, verified sender identity.

- Use the "Outreach Email Configuration" section under the General tab in the Admin Panel to enable or disable agent email sending for the entire organization, with a clear indication that when disabled, no project in the organization can send agent email mentioned in its instruction.
- Select the preferred Email Provider, choosing between an SMTP server or SendGrid (Web API), so that email delivery aligns with the organization's existing mail infrastructure. The configuration fields update automatically to match the selected provider.
- When SMTP server is selected, define the full SMTP configuration, including SMTP Server Address, Port, Use SSL/TLS, From Address, From Display Name, SMTP Username, and SMTP Password, ensuring secure and correctly attributed delivery.
- When SendGrid (Web API) is selected, provide the SendGrid API Key along with the From Email (verified sender) and From Display Name, so that email is sent through SendGrid from an approved, verified sender. The API Key field can be left unchanged to retain the current key when updating other settings.
- Validate the configuration before going live by using "Send Test Email" under Test Email button to confirm that credentials and connectivity are working as expected, regardless of the selected provider.
- Create and maintain reusable email content through "Manage Templates," so that agent notifications remain consistent and on-brand across the organization.
- Download a complete Email Audit Log as a CSV file for traceability, compliance, and troubleshooting of all agent-sent email.

c. Daily Credit Consumption Alerts

Administrators can now monitor and govern daily credit usage by setting warning thresholds, designating notification recipients, and enabling automated safeguards. This helps organizations stay ahead of unexpected consumption and protect the monthly quota from being exhausted by a small number of high-usage agents.

- Use the "Daily Credit Consumption Alerts" section under the General tab in the Admin Panel to send a warning email whenever the organization's daily credit utilization exceeds a selected percentage.
- Set a Warning Threshold as a percentage of the monthly quota, so that alerts are triggered at a level appropriate to the organization's usage patterns.
- Specify one or more Notification Email Addresses to ensure the right stakeholders are informed as consumption approaches critical levels.
- Enable Agents4DevOps to automatically disable a defined number of the highest-consuming agent(s) once today's credit utilization reaches a selected percentage of the monthly quota, providing an automated safeguard against runaway usage.
- Configure both the count of top consuming agents to disable and the utilization percentage that triggers the action, giving administrators fine-grained control over the automated response.

d. Agents4DevOps Recommended Models

Administrators can now configure which AI models power each Agents4DevOps module by defining a prioritized list of models per module. This gives organizations centralized control over model selection and allows lower-ranked models to serve as fallbacks, ensuring consistent behavior across all modules.

- Use the "Agents4DevOps Recommended Models" section under the General tab in the Admin Panel to configure AI models for each Agents4DevOps module from a single, structured table of Module Name and Ranked Models.
- Assign one or more ranked models to each supported module, including AgentCodeGenerationModel, AgentExecutionModel, AgentInstructionGenerationModel, AgentRegenerationModel, ChatAssistant, SkillGenerationModel, and SkillRegenerationModel.
- Order the models within each module to establish priority, so that the system uses the highest-ranked model first and falls back to subsequent models as needed without disrupting agent operations.
- Add or remove models per module directly from the table, with each selected model displayed as a removable entry for quick adjustment.

e. Configure Chat Assistant Model

Administrators can now select which AI model powers the Chat Assistant across the entire platform. This ensures a single, consistent conversational experience for all users, while remaining aligned with the models approved in the organization's recommended models configuration.

- Use the "Configure Chat Assistant Model" section under the General tab in the Admin Panel to configure the AI model used for the Chat Assistant across the platform.
- Select the desired Chat Assistant model from a dropdown of approved options, applying the choice consistently for all users.
- Choose only from models that are configured in the Agents4DevOps Recommended Models section, ensuring that Chat Assistant selection stays governed by the organization's central model configuration.

f. Global Tool Call

A new Global Tool Call section in the Admin Panel gives administrators centralized control over which tool call are available to agents. This allows organizations to govern behavior, reduce risk, and enable only approved capabilities.

- Manage all tool calls from a structured, collapsible list under the General tab, grouped by tool and section, with each individual tool call displayed as "Name & Description" in a clear, readable format.
- Toggle entire tools section on or off, use section-level "Enable All / Disable All" actions, or control individual tool call via per-row toggles, with all settings persisted using the existing configuration pattern.
- Start from a secure, permissive default where all skills are enabled on first load, and then selectively disable capabilities that should not be used by agents, with runtime enforcement that blocks disabled tool call and surfaces clear messages indicating that a tool call is disabled in Global Tool Call settings.

g. Global and Project-Level Permission Management

The release introduces a comprehensive permission model that let administrators control what actions users can perform with agents and jobs, both globally and per project. This

ensures that sensitive operations such as executing agents, deleting jobs, or managing tasks are only available to authorized roles.

- At the organization level, use the “Global Permissions” subsection in the Admin Panel to configure twelve key actions including: Manage Agent (it controls Add, Edit, and Clone action of Agent), Delete Agent, Execute Agent, Enable Agent, Customize Agent, Manage Skill (it controls Add, Edit, and Clone action of Agent), Delete Skill, Enable Skill, Customize Skill, Job Explorer, Delete Job, and Open Task, each with a simple Allow/Deny option.
- At the project level, open Project Settings → Agents4DevOps to see all teams and core Azure DevOps groups (Project Administrators, Build Administrators, Project Valid Users, Readers, Contributors) and configure the same twelve actions per team or group.
- For each action at project level, choose Allow, Deny, or Not set, with clear indication when a value is inherited from global settings (e.g., “Allow (inherited)” or “Deny (inherited)”), and the ability to revert to inheritance by selecting Not set.
- When an action is denied, the corresponding UI elements are disabled consistently.
- Ensure that permissions for Job Explorer, Delete Job, and Open Task are applied consistently across both the master Jobs tab and agent-specific job views, providing predictable behavior for users regardless of where they access job information.