

Prompt to Researcher:

"I am writing a research paper on the new agents by Microsoft, the "Researcher" and the "Analyst". Please help me find facts about these agents. In specific, I would like to know how they are architected, what data models they are using, any technology used from organizations like OpenAI, and Microsoft technology. I would like to understand the various use cases for these specific agents, any specific limitations or features or functions that end users need to be aware of. I would also like to know customer and user feedback. Please provide any insight into market packaging, positioning and availability when it comes to languages, geographical roll out, and so on. Use any sources you can find on the internet, including YouTube videos or articles from the community as well as published articles by Microsoft or other vendors, research institutions or media. I would also like to understand how these agents relate to other similar types of AI-services from other vendors."

Researcher's follow up questions:

"I see you want me to help you gather detailed information about the new Microsoft agents, "Researcher" and "Analyst," including their architecture, data models, technology, use cases, limitations, user feedback, market positioning, and comparisons with similar AI services from other vendors.

To make sure I understand your request:

1. Are there any specific sources or types of sources you prefer for this information (e.g., official Microsoft documentation, community articles, YouTube videos)?
2. Do you have a preferred format for the final report (e.g., a detailed written report, a presentation, a summary)?
3. Are there any particular geographical regions or languages you are interested in regarding the availability and rollout of these agents?

If you'd like me to proceed, tell me to "go ahead" and I'll use my best judgement."

My answer:

"Since this is such a new technology, you can go very broad in your research when it comes to information sources. Make sure to reference your sources. When it comes to the format of the report, you can just write it in a text format (I will use it as input for my article) when it comes to geographical overview primarily US, Europe and the Middle East (but with that said, a "global" perspective is preferred)."

Microsoft 365 Copilot “Researcher” and “Analyst” Agents: In-Depth Analysis

Researcher: AI-Powered Research Assistant

Helps users conduct deep, multi-step research by combing through work data (emails, files, chats, etc.) and the web to produce detailed reports with sources.

Analyst: AI-Powered Data Analyst

Acts as a virtual data scientist that analyzes raw data (spreadsheets, databases) step-by-step, running code as needed to find trends, generate visualizations, and deliver insights.

Overview

Microsoft introduced **Researcher** and **Analyst** in early 2025 as two first-of-their-kind AI *reasoning agents* integrated into the Microsoft 365 Copilot ecosystem^[1]. These agents are designed to function like on-demand experts embedded in the tools people already use (such as Word, Excel, Teams), augmenting the standard Copilot experience with deeper analytical capabilities^[2].

Researcher is positioned as an assistant for **deep, multi-step research** tasks, capable of securely fusing internal enterprise information with external web data to produce comprehensive, well-sourced reports^{[1][1]}. **Analyst** is aimed at **advanced data analysis**, functioning as a “virtual data scientist” that can take raw data (e.g. Excel sheets, CSV files, databases) and turn it into actionable insights through iterative reasoning and computation^{[1][2]}.

Together, these agents represent Microsoft’s push to provide “*expertise on demand*” within the flow of everyday work^[1]. In effect, Microsoft 365 Copilot is evolving from a single general assistant into a **team of AI specialists** – with Researcher tackling complex research inquiries and Analyst handling data-heavy analytic problems^{[3][3]}. Both agents became generally available to all Microsoft 365 Copilot licensed users by June 2025, after an initial preview via the “Frontier” early access program^{[4][1]}.

Architecture and Design

Microsoft's Researcher and Analyst agents are built on a foundation of **advanced large language models combined with tool integration and multi-step reasoning techniques**. They do more than just answer questions in one go; under the hood, they break down tasks, invoke data search or code tools, and iteratively refine their approach. This section examines how each agent is architected and the data models and technologies they use.

Researcher Agent Architecture

The Researcher agent employs a **structured, multi-phase reasoning process** that mimics how a human researcher would approach a complex question[5][5]. When a user poses a query, Researcher first enters a **planning phase (P_0)** where it analyzes the query and determines a game plan – often asking the user clarifying questions if necessary to sharpen the scope[5][5]. Once the plan is set, the agent enters an **iterative research loop** consisting of three repeating steps (often summarized as *Reason* → *Retrieve* → *Review*)[5][5]:

- **Reasoning (R_i):** The agent decides which sub-problem or aspect to tackle next and formulates hypotheses or what information is needed[5].
- **Retrieval (T_i):** It then searches for relevant information to fill that knowledge gap – this includes querying enterprise data (SharePoint documents, emails, Teams chats, meeting transcripts, etc.) as well as performing web searches[5][5]. Through Microsoft Graph connectors, Researcher can also pull content from third-party business systems like Salesforce, ServiceNow, or Confluence[1][6], ensuring both internal and external sources are tapped.
- **Review (V_i):** The agent evaluates the retrieved material, determines its relevance to the query, and logs any new “insights” gained into a working memory (a “scratch pad”)[5][5].

This loop repeats, with each iteration adding to the cumulative insight pool, until the agent detects **diminishing returns** – i.e. new iterations yield minimal new information[5][5]. At that point, Researcher moves to a **synthesis phase**, where it compiles all the gathered findings (denoted as I_m) and generates a coherent, structured report as the final output[5][5]. The report includes explanations, key insights, and citations for sources used, providing traceability for each fact[5][3]. This approach of iterative deep reasoning means Researcher may take longer to respond than a standard single-turn Copilot chat (it “thinks” more), but in return it delivers a far more thorough and accurate result[5].

Underlying Model: Researcher is powered by an advanced large language model specialized for research tasks. Microsoft confirms that its implementation **leverages OpenAI's upcoming "deep research" model**, described as a version of OpenAI's O3 model family that has been **fine-tuned for complex research and reasoning**[\[5\]](#). This is an evolution of the GPT-4 class of models – one that particularly excels at multi-step inference and handling large volumes of text. Benchmarks illustrate the model's prowess: for example, it achieved ~26.6% accuracy on the challenging *Humanity's Last Exam (HLEx)* test and ~72.6% on the GAIA reasoning benchmark, outperforming prior approaches[\[5\]\[5\]](#). The Researcher agent's model is augmented with what Microsoft calls the **Copilot tools/Control system**, an orchestration layer that grants the model the ability to perform actions like searching documentation, opening files, and extracting content[\[5\]\[5\]](#). During its reasoning process, the agent can "open" a document or "find" text within it via these tools, using an interface similar to how it was trained to use a web browser or other tools[\[5\]](#). This design helps the AI interact with enterprise data sources in a controlled, auditable manner. Microsoft has also integrated the **enterprise knowledge graph and user context** into Researcher's reasoning. This means the agent is aware of organizational specifics – for instance, it knows project names, team hierarchies, and domain-specific terminology from your tenant's data, which allows it to personalize its search and better interpret queries in your business context[\[5\]\[5\]](#). (E.g., if a user asks about "Project Olympus," Researcher can discern whether that refers to an internal project or an external concept, and ask nuanced follow-ups accordingly[\[5\]](#).) This personalization mitigates the model's lack of familiarity with company jargon and yields more relevant results.

In summary, Researcher's architecture marries a **powerful, research-optimized LLM** with a sophisticated retrieval framework. It performs broad but shallow searches across heterogeneous sources and then zooms in on the most pertinent domains and documents as it iterates[\[5\]\[5\]](#). It can even hand off certain subtasks to other specialized AI agents if needed – for example, invoking a "Sales data" agent for a time-series forecast within a larger research question[\[5\]\[5\]](#). This modular, tool-using, multi-step design is why Researcher can tackle complex questions that stump simpler one-shot AI assistants, albeit with some trade-off in speed.

Analyst Agent Architecture

The Analyst agent is architected to tackle **structured data analysis** problems by employing an iterative, code-assisted reasoning process. It uses a **chain-of-thought (CoT) approach** similar in spirit to Researcher's, but optimized for data manipulation and quantitative reasoning[\[7\]](#). Instead

of focusing on search across text sources, Analyst focuses on **computational steps**: it will methodically parse provided datasets, generate code to transform or analyze data, and adapt based on intermediate results[\[7\]\[7\]](#).

At the core of Analyst is an “**advanced reasoning model**” **post-trained on OpenAI’s O3-mini** model, specifically tuned for analytical tasks[\[7\]](#). OpenAI’s *O3-mini* can be thought of as a compact but powerful language model variant (part of the same family as GPT-4) that Microsoft further reinforced with domain-specific training. Microsoft describes Analyst as a “*reasoning-powered agent*” that serves as a “*virtual data scientist*”, given its ability to not only converse but also to **generate and execute code on the fly** as part of its reasoning[\[7\]\[7\]](#).

Here's how Analyst’s problem-solving loop works in practice:

- The user begins by posing a question about data and usually attaching one or more data files or selecting data sources (e.g., Excel workbooks, CSVs, database tables, etc.) to analyze[\[8\]\[8\]](#). (Analyst can also work with cloud files from OneDrive/SharePoint if given access, or the user can upload files – a step akin to providing the “sandbox” with data to crunch.)
- **Iterative Analysis:** Upon receiving the query and data, the Analyst agent breaks the task into smaller steps internally. It might first inspect the structure of the files, then attempt initial computations. Crucially, Analyst’s model can **write Python code** (or other script) to perform calculations or data transformations during each reasoning step[\[7\]\[7\]](#). This code-execution ability is a defining feature: the agent effectively has a built-in data analysis Jupyter-style environment. For example, it may generate Python code to clean a dataset (fixing a delimiter issue or filling missing values), then run that code, evaluate the output, and decide the next step[\[7\]\[7\]](#).
- The chain-of-thought architecture means Analyst doesn’t jump straight to final conclusions; it **hypothesizes, tests, and refines iteratively**[\[7\]](#). At each stage, it checks if the result makes sense or if further analysis is needed, adjusting course if an error occurs. Notably, if errors happen (say a piece of code fails), the agent can debug and try again, exhibiting a form of automated error recovery that is uncommon in standard LLM responses[\[7\]\[7\]](#). Microsoft specifically notes that Analyst “*takes as many steps as necessary*” and can **gracefully recover from mistakes**, mimicking how a human analyst would troubleshoot and persist[\[7\]\[7\]](#).

Under the hood, **reinforcement learning (RL)** techniques were key to developing this agent. Microsoft took the base OpenAI O3-mini model and *post-trained* it through extensive RL with a focus on stepwise reasoning and coding skill[\[7\]](#). They defined rule-based reward functions to encourage the AI to pursue longer reasoning paths when needed, perform incremental information gathering, and successfully execute code to solve problems[\[7\]\[7\]](#). The training

encompassed a wide variety of real-world enterprise scenarios: the team built datasets spanning Excel, CSV, JSON, XML files, SQL databases, Power BI reports, etc., with tasks ranging from simple computations to complex multi-file analyses[7][7]. By exposing the model to messy “real” data (like misformatted files or hidden cells in spreadsheets), they conditioned Analyst to handle the kind of data issues that often trip up normal LLMs[7][7]. For instance, the official blog recounts a case where Analyst had to pull a critical number buried mid-sheet in an Excel file and deal with a TSV file that was incorrectly using commas – tasks it managed by identifying sheet names and fixing the delimiter before proceeding[7][7].

As a result of this design, Analyst can do things like: explore data tables to find relevant fields, generate summary statistics, join data from multiple sources, identify outliers, and even create charts or visualizations – all through a reasoning process that **combines natural language understanding with programmatic data analysis**. The agent outputs its findings in an “**analysis report**” (often with charts or tables included when appropriate) written in plain language explaining the results[8][8]. Importantly, it also **exposes the code it used** along the way. In the interface, users are able to view the Python code that Analyst ran for each step, lending transparency to the analysis[1][2]. This builds trust, as an end-user or data professional can inspect or even copy that code to verify calculations independently.

Key Technologies Used: Analyst’s heavy use of Python for computation is enabled by a **secure sandbox environment** within Microsoft 365 Copilot. It’s analogous to OpenAI’s “Code Interpreter” concept – the agent can execute code in a controlled, ephemeral runtime that has access to the provided data but is isolated for security. Microsoft indicates the model “writes, executes, and verifies Python code within a controlled execution environment” as part of its reasoning cycle[7]. This means, for example, if you ask Analyst to forecast trends from sales data, it might write a short script using a forecasting library or custom logic, run it, and then interpret the output to explain the trend. All this happens behind the scenes in seconds or minutes, depending on complexity.

Similar to Researcher, Analyst also relies on the synergy of OpenAI’s model and Microsoft’s own systems. Microsoft’s **orchestration** layer feeds the model the necessary tools (like the code execution tool) and data attachments. The **Copilot Control System** ensures any data used stays within compliance and that the agent respects the user’s permissions to data sources[1][1]. From a model standpoint, OpenAI’s *O3-mini* is acknowledged as providing strong “reasoning and coding” abilities with high efficiency – Microsoft then built on it with additional training for enterprise use

cases. In January 2025, Microsoft even made the O3-mini model available in the Azure OpenAI Service for developers, highlighting its role in enabling advanced reasoning scenarios.

In summary, Analyst's architecture is about **combining language, reasoning, and code execution**.

It uses a fine-tuned OpenAI model that engages in chain-of-thought reasoning, with the unique ability to inject code at each step to handle calculations or data manipulation. This design allows Analyst to tackle problems like, *"Compare sales performance across regions and quarters and highlight key trends,"* where it might need to aggregate figures from multiple files and generate a chart^[8] – tasks that purely language-based AI would struggle with or do incorrectly. Analyst, however, can literally calculate the answer by writing code, leading to more accurate and thorough analysis results.

Use Cases and Capabilities

Microsoft's Researcher and Analyst agents are intended to address different categories of workplace scenarios. Below, we outline the primary use cases for each agent and what capabilities they bring to end-users.

Researcher Agent – Use Cases

Researcher is best at tackling scenarios that involve **information synthesis, strategic research, and report generation**. It shines when a user needs to gather insights from a wide array of sources and formulate a coherent analysis. Some key use cases include:

- **In-Depth Research Reports:** Compiling comprehensive reports on a topic by pulling together internal knowledge and external intelligence. For example, a product manager can ask Researcher to *"build a detailed go-to-market strategy based on all our project data and current market trends"*. The agent will gather context from internal documents (plans, meeting notes, sales data) and relevant web articles or industry reports, resulting in a thorough plan with both internal metrics and competitive analysis included^[1]. This could cover market sizing, competitor offerings, SWOT analysis, etc., in one package.
- **Competitive and Market Analysis:** Researcher can investigate a market or competitor by reading through news, analyst reports, and internal sales materials. Microsoft gave an example that you could *"identify whitespace opportunities for a new product based on emerging trends and internal data"*^[1]. In practice, you might prompt: *"Find any emerging customer needs in the renewable energy market and see if our product line has gaps there."* Researcher would search internal R&D docs, customer feedback emails, and external

market trend reports to highlight potential opportunities that competitors aren't covering, complete with supporting facts.

- **Client Preparation and Dossiers:** The agent can assemble briefs for client meetings or sales pitches by combining internal data on the client/account with industry news. For instance, *"Prepare a quarterly business review report for Client X, including work we've done to date and the latest industry developments relevant to them"* is a task ideally suited for Researcher[1]. It will summarize project milestones from your files and chats, and merge that with any new regulations or market changes affecting the client's domain, yielding a ready-to-present briefing document.
- **Policy or Strategy Papers:** Internal strategy teams can use Researcher to draft documents like policy briefs, investment rationales, or research whitepapers. Because it can cite sources, the output is easy to verify. E.g., a user could ask, *"Analyze how recent tariff changes might impact our supply chain and suggest mitigation strategies,"* and Researcher would scour internal supply chain reports plus external news about the tariffs to provide an informed analysis (Microsoft indeed suggests it for assessing impacts such as new tariffs on business lines)[3].
- **General Knowledge Queries with Context:** It can answer complex questions that aren't answerable from a single source. For example, *"Give me a summary of how AI is used in healthcare, focusing on any projects our company has done in that space,"* would prompt it to combine public info (AI in healthcare trends) with private info (maybe an internal project on a healthcare AI pilot) into a single narrative. This goes beyond a web search result by integrating the user's context securely.

Capabilities of Researcher: To accomplish the above, Researcher brings a few notable capabilities:

- It performs **multi-source information gathering** across both **internal** repositories and the **web** in one query[3][3]. This means less manual hunting for the user. For example, if investigating a topic like "blockchain in finance," Researcher might pull relevant emails from your finance team, SharePoint files of past research, and recent web articles from financial outlets, all together.
- It provides **smart summaries**. Rather than just dumping raw text, it synthesizes and **summarizes findings in plain language**, focusing on key points and insights[3][3]. The output is a structured report rather than a loose collection of excerpts.
- It can **identify trends, patterns, or insights** across the collected info. For example, it might notice that multiple sources mention a particular risk factor and call that out. In Microsoft's testing, if Researcher sees a common thread (say, feedback about a product feature from different sources), it will highlight that as a noteworthy insight[3][3].
- Researcher engages in **interactive refinement**. If a user's question is broad or ambiguous, the agent often asks follow-up questions to clarify requirements or scope (e.g., "Do you

want to focus on a specific region?”)[5][3]. This back-and-forth helps it zero in on exactly what the user needs. Users can answer these or skip them, but providing guidance can greatly improve relevance.

- Crucially, it outputs **fully cited content**. Every factual statement comes with a citation or reference to the source (much like an academic paper), so the user can verify where information came from[5][3]. This transparency is critical in a work setting; you can even ask Researcher to only include information from authoritative or specific sources if needed for confidence[3].

In essence, Researcher is like a tireless research analyst on your team: it will trawl through *both your company's knowledge bases and the vastness of the internet*, then condense everything into a digestible report. Microsoft's CVP Jared Spataro summarized it as “*an agent that can analyze vast amounts of information with secure, compliant access to your work data and the web to deliver expert insights on demand*”[3]. Early use cases from customers included things like market entry studies, competitive landscape reports, and client-ready insights that previously might have taken analysts days of work to assemble.

Analyst Agent – Use Cases

Analyst is targeted at scenarios involving **data analysis, business intelligence, and forecasting** – especially for users who are not expert analysts or who want to automate the grunt work of parsing data. Key use cases for Analyst include:

- **Exploratory Data Analysis:** Given one or multiple datasets, Analyst can quickly summarize what's in the data, calculate relevant statistics, and present findings. For example, a small business owner can upload a few CSV files (sales data, marketing leads, customer feedback) and ask, “*What factors are driving our sales growth this quarter?*” Analyst might compute correlations or trends across those files and report something like “*Sales are up 10% in Q2, with the new product line contributing 80% of growth; increased website traffic (up 50%) is correlating with sales spikes, suggesting marketing campaigns are effective*”, with charts of sales by product and traffic vs. sales[8][8].
- **Generating Reports from Raw Data:** Users can obtain **ready-made reports and visuals** without manually building pivot tables or charts. For instance, “*Compare sales performance across regions and quarters and highlight key trends*” is a query that Analyst can handle by examining the provided regional sales Excel files and producing a concise report with a table or chart of regions vs quarters, plus narrative insight on which regions outperformed and any seasonality[8]. The agent's ability to create easy-to-read summaries with visuals means it can effectively automate a task that might take an analyst hours in Excel.

- **Trend Analysis and Outlier Detection:** Analyst excels at finding patterns in data. A user might ask, *"Identify any anomalies or outliers in our finance spreadsheet"* or *"Show me trends in customer engagement over the past year."* The agent can run the necessary computations (like computing statistical outliers, trends over time, etc.) and then explain, for example, *"Q4 had an unusually high expense in marketing (3σ above mean), which stands out as an outlier"* or *"Customer engagement has been steadily rising 5% month-over-month with a notable spike in July due to campaign X"*. It effectively does the heavy lifting of scanning thousands of rows to flag anything noteworthy[8][8].
- **Forecasting and "What-If" Analysis:** Thanks to its coded reasoning, Analyst can even perform basic predictive analysis. Use cases like demand forecasting or financial projections are within scope. A user could request, *"Project our revenue for next quarter based on the last 2 years of data"*, and Analyst might use a time-series model in Python, then report *"Expected Q4 revenue is \$1.2M ($\pm 0.1M$), which is a 8% increase YoY, assuming current growth rate holds"*, possibly providing a chart of the projected trend[1]. Similarly, *"If we increase pricing by 5%, how might that affect our monthly sales?"* could prompt the agent to run a scenario analysis (if the data to infer price elasticity is available).
- **Customer and Marketing Insights:** Non-technical users in sales or marketing can leverage Analyst to interpret data like customer surveys, CRM exports, or campaign metrics. For example, *"Analyze customer behavior data and show any interesting segments"* might lead the agent to cluster or segment customers (again via behind-the-scenes code) and then describe *"Three key customer segments identified: (1) Repeat high-value buyers (10% of customers, 30% of revenue), (2) One-time bargain shoppers... etc."* along with suggestions on how to target each group. Microsoft specifically notes Analyst can create *"customer behavior visualizations"*[1] or highlight top-performing products/services based on the metrics provided[8].
- **Routine Reporting Automation:** Many businesses have regular reports (weekly sales reports, monthly financial summaries). Analyst agent can potentially automate these. A finance team member could ask, *"Summarize this month's financial report and highlight any variances from last month"* after attaching the monthly Excel file. Analyst would calculate differences, produce a summary (e.g., *"Revenue grew 5%, expenses up 2%, notable increase in travel costs..."*), and even generate charts for the major line items. This frees users from manual spreadsheet work each period.

Capabilities of Analyst: To serve these cases, Analyst offers:

- **Natural language querying of data:** Users don't need to write formulas or code; they can ask questions in plain English (or other supported languages) about their data[8]. Analyst figures out the analytical steps needed.
- **Multi-dataset analysis:** Analyst can merge and analyze multiple files together. Users can attach more than one source (which it then can join or cross-reference as needed)[8]. For example, combining an Excel sales report with a CSV of website analytics in one question –

the agent can reconcile them if there's a common field (like date or product ID) and draw insights across both.

- **Automated calculations and statistics:** It will compute things like sums, averages, growth rates, correlations, etc., on the fly. A prompt like *"What is the average customer spend and who are the top 5 customers?"* would lead it to calculate those from the data. The user does not have to manually do any math or Excel work – Analyst handles it, then gives the results in narrative or table form[\[8\]\[8\]](#).
- **Charts and visualizations:** The agent can generate simple charts (bar charts, line graphs, etc.) to include in its answer for better understanding[\[8\]](#). If you ask for trends or comparisons, the result often comes with a chart image. This turns raw numbers into visual insights automatically.
- **Recommendations and next steps:** Analyst doesn't just report the data; it often will include a brief interpretation or suggestion. For instance, after spotting a trend, it might add *"Consider increasing stock for Product A in Region North, as it's consistently outperforming others"*. This prescriptive element helps non-experts know what to do with the data findings[\[9\]\[9\]](#).
- **No code or expertise required:** The whole idea is to empower people who are not data scientists. Analyst *"is like having a skilled data analyst on hand"* that lets you avoid writing complex formulas or scripts yourself[\[8\]](#). This lowers the barrier for advanced data analysis—tasks like regression or parsing JSON data can be done by the AI whereas before only those with technical skills could do them easily.
- **Transparency via code view:** For power users or skeptics, the ability to inspect the Python code that was run is extremely valuable[\[1\]\[2\]](#). One can see exactly how a figure was calculated or which algorithm was used for a prediction. This feature sets it apart from typical "black box" AI: if something looks off, a data analyst in the team can review the code and adjust the approach if needed.

In short, Analyst is particularly useful for **making sense of complex or large datasets quickly**. It turns the often tedious process of cleaning data, running analyses, and creating charts into an interactive Q&A experience. This can benefit roles like business analysts, sales ops, marketing, or finance professionals who have the data and questions, but not enough time or perhaps skills to manually do deep analysis each time. As Microsoft's demonstration showed, it can reduce tasks that might take days (like consolidating results from many Excel sheets and drawing insights) into minutes[\[6\]\[6\]](#). A practical small-business example given is using Analyst to combine inventory, sales, and customer review data to determine which products are performing best and why[\[8\]](#) – something that would normally require hours of sorting and chart-making, but the agent can produce an answer with a couple of file attachments and a single question.

Key Features and User Experience

From an end-user's perspective, interacting with the Researcher and Analyst agents is meant to be straightforward. They are accessible through the **Microsoft 365 Copilot interface**, which can be reached via [Microsoft365.com](https://microsoft365.com) or within Office applications. In the Copilot web interface, these agents appear under an "Agents" section (often via an "All Agents" or "Built by Microsoft" catalog)[\[10\]](#). A user can select **Researcher** or **Analyst** from the list of available Copilot agents to start using them[\[10\]](#).

Once opened, the **workflow is conversational**: you pose a prompt or question in natural language, and the agent will process it and eventually return a result (which could be a multi-page report or a set of analysis findings, depending on the agent). Some notable features of the user experience include:

- **Clarifying Questions:** Especially with Researcher, after you submit a query the agent may ask you one or more follow-up questions if it's unsure what exactly you need[\[5\]\[11\]](#). For example, if you ask a broad question like "Analyze the impact of AI on education," Researcher might respond with something like, "Do you want to focus on higher education or K-12, or cover both? Any specific region or timeframe?" This interactive dialogue helps narrow down the scope. You can answer to refine the direction or simply tell it to proceed if you meant it broadly[\[11\]](#). This makes the agent feel somewhat collaborative, as it tries to ensure it's on the right track before diving into a lengthy analysis.
- **Attaching Data (for Analyst):** For the Analyst agent, the interface explicitly allows you to attach files or select data sources when you ask a question[\[8\]\[8\]](#). There's typically an "Attach content" button (with an upload icon or OneDrive icon) where you can pick one or multiple datasets for the agent to analyze. You can search your OneDrive/SharePoint directly from this prompt (for example, typing part of a file name to find an Excel file in your cloud storage) or upload files from your device[\[8\]\[8\]](#). This step is important – Analyst will only analyze data you provide or point it to. In contrast, Researcher by default will search across your organizational data (with your permissions) without manual attachments; it has enterprise search capabilities given appropriate access.
- **Progress and Transparency:** When the agents are working on a request, you might see indications of progress. In some demos, Microsoft has shown that Researcher will list out sub-tasks it's working on or sources it's retrieving (almost like a status update), though in the current product the UI might simply show a loading animation with messages like "Analyzing data...". Analyst, as it runs code, might show the code pane updating in real-time, which can be watched by the user. This transparency – seeing code execute or seeing that it's fetching certain documents – is unique. It addresses the typical black-box nature of AI. In Analyst's case, the user can expand a panel to **view the Python code** that the agent is generating and running during the session[\[2\]\[1\]](#). This is read-only (the user doesn't need

to write anything), but it's informative for those curious or for validating what the agent did. For Researcher, after it produces the answer, you will see **in-text citations (like reference numbers)** that link to the source documents or webpages it used^[3]. A user can click these to open the source material, ensuring every claim can be traced.

- **Structured Outputs:** Both agents strive to present information in a **clean, organized format** rather than a single long paragraph. Researcher's output is often a report with section headings, bullet points for key findings, and enumerated references^{[11][11]}. It may also include a list of "Next steps" or recommendations at the end of a report^[11]. Analyst's output typically includes text explanations and, when relevant, inserted charts, tables, or lists. For example, after analyzing data it might show a bar chart of sales by region next to the text explanation in its chat answer. The idea is that the result is **ready to use** – you can copy it into a Word document or PowerPoint with minimal editing. In fact, Microsoft intends that you should be able to directly take a Researcher report and share it or "*edit it in Microsoft apps like Word or Pages*" as a starting point for a deliverable^{[11][11]}.
- **Citations and Source Drill-down:** As mentioned, Researcher heavily emphasizes source citing. Users can click on a citation to see the portion of the source document or webpage that backs up a statement. This fosters trust (and is also useful for getting more detail, as the report might summarize something that you then want to read about in full context). Analyst, while citing code and data steps, deals with facts derived from user-provided data, so citation in the traditional sense (like a document link) is not applicable unless it uses external data. However, if Analyst pulls in data from a database or a file, it will indicate what source (file name) it used for each figure.
- **User Guidance and Control:** Users remain in control of the agents. You can stop an analysis partway if it's taking too long, or instruct the agent with additional prompts. For instance, if Researcher's answer is not exactly what you need, you can ask follow-up questions or say "focus more on X aspect" to iterate further. With Analyst, if you realize you gave it the wrong file, you can correct that and re-run. Microsoft is also adding controls in the admin center so that organizations can govern what data sources agents can or cannot use (more on that later in the limitations). From the user side, one can also tell Researcher to *exclude* certain sources or *only use internal data*, etc., by specifying that in the prompt.

Overall, the design goal is that using these agents feels like interacting with a very smart assistant: you ask in normal language, the assistant might ask a clarifying question or two, then it shows you an answer complete with all the evidence and even the workings (code) it used. This is quite novel compared to earlier AI assistants that would just give a final answer without context. End-users do not need to know anything about AI or coding to benefit – they just need to know how to ask for what they want (just as they would ask a human colleague). If they can formulate the question or the task, the agent will attempt to figure out the rest.

One thing to note: because Researcher and Analyst operate within the Microsoft 365 Copilot app framework, they inherit the same enterprise security, compliance, and privacy protections as the rest of Copilot. All data stays within the tenant boundaries when being processed (with cloud AI models handling it under Microsoft's enterprise agreements). Microsoft's Copilot Control System ensures, for example, that Researcher will only retrieve documents the user has permission to access, and all outputs respect compliance settings (no leaking of confidential info to unauthorized users)[\[1\]](#). This means from the user perspective, they can trust the agent to handle sensitive internal information as securely as other M365 services do, which is critical for adoption in corporate environments.

Limitations and Considerations

While Researcher and Analyst are powerful, they do come with **certain limitations and usage considerations** that end-users and organizations should be aware of. As cutting-edge AI agents, they are not without constraints in performance, availability, and accuracy. Here are some key limitations and cautions:

- **Usage Limits:** Microsoft has implemented **quota limits** on these agents (at least as of general availability in 2025) to manage load and ensure quality of service. Specifically, the two pre-built agents share a pool of **25 queries per user per month** – that is, each licensed user can run up to 25 combined Researcher+Analyst requests in a month[\[4\]](#). It's not 25 each; if you use Researcher 15 times and Analyst 10 times, that hits the limit. This cap was introduced likely because these agents are resource-intensive (each query might involve many sub-queries, API calls, or code executions under the hood). For comparison, the regular Copilot chat for simpler tasks doesn't have such tight limits. Heavy users will need to plan their agent queries or request increased limits if Microsoft offers that in the future. (Notably, during the initial preview in Frontier, Microsoft had stated "no limits on usage... subject to change" while in development[\[10\]\[10\]](#) – clearly, they decided to impose limits for GA.)
- **Latency and Speed:** Both agents perform multi-step processing, which means **responses can take longer** than a normal AI chatbot answer. Researcher, for example, might spend several minutes on a complex query as it methodically gathers info. Users have reported cases where Researcher took quite a long time (e.g. 20-30 minutes) and in some cases timed out without completing[\[12\]\[12\]](#). Microsoft in documentation has acknowledged that these "deep reasoning" workflows are slower than the near-instant answers people might be used to from Copilot Chat[\[5\]](#). The expectation is that if you're asking for a comprehensive report, you might start the agent and wait a bit. There isn't a strict SLA on how fast it returns; it depends on the complexity of the task and the amount of data.

Analyst, dealing with code and data, may also take some time especially on large datasets or complicated computations. The user interface will show it “working,” but patience is needed for the best results. Organizations should be aware that these are not real-time instantaneous responses, so they’re suited for when a user can wait a minute or two for a richer answer.

- **Complexity Limits:** Although these agents are designed for complex tasks, there **are** still boundaries on what they can handle. Extremely broad or overly ambitious prompts might fail. For example, one user’s attempt to generate a very elaborate organizational analysis (involving multiple sub-topics and a provided Power BI dashboard) overwhelmed the Researcher agent, resulting in an error with no useful output[\[12\]\[12\]](#). The agent effectively gave up (“Sorry, something went wrong”) instead of partially answering. This suggests that there are practical limits to the complexity – perhaps in terms of how many iterations the agent will attempt or how much data it can juggle at once. If a request demands too many divergent tasks (e.g. “Give me everything about X, plus a full SWOT analysis, plus an implementation plan”), the system might not yet gracefully scale to that. Users might need to break very large requests into smaller questions for now.
- **Graceful Degradation:** Building on the above point, one identified limitation is that the agents may not always **gracefully degrade** when they can’t fully solve something[\[12\]\[12\]](#). Ideally, an AI agent would at least give a partial answer or explain which part of the task was too difficult. In the current version, however, a failure can result in a generic error with no insight into which step failed. In the example case, the user got no partial results or explanation[\[12\]](#). This is an area for improvement – future updates might handle partial fulfillment better or provide error details (“I couldn’t access the data source” vs “the question is too broad”), but at GA it’s somewhat all-or-nothing.
- **Language Support:** At launch, there is a discrepancy in language coverage between the two agents. **Researcher supports 37 languages** (essentially a broad global coverage of major languages), whereas **Analyst supports only 8 languages** initially[\[4\]](#). The supported languages are not listed in the source, but likely Analyst’s initial languages include English, French, German, Spanish, Japanese, Chinese and a few others, whereas Researcher’s 37 would cover most European and several Asian languages, possibly Arabic and others. Microsoft has indicated more language support for Analyst is “coming soon”[\[4\]](#). The reason for the gap is likely that the data analysis model and its code generation needed extra tuning per language. So, if a user tries to converse with Analyst in an unsupported language at GA, results may be poor or it may fallback to English. Organizations in non-English locales should check which agents are fully functional in their language. (During the preview phase, both agents were available **only in English** as they were still under development[\[10\]](#), so the GA expanded to many languages for Researcher and a subset for Analyst.)
- **Data Attachment Limits:** For Analyst, although it can handle multiple files, there are practical limits (not explicitly stated publicly) on **file size or number of files** it can process

in one go. If you feed it huge datasets (say a million-row spreadsheet), it might not be able to load that entirely or it could hit execution time limits. Similarly, extremely large images or JSON files etc., could be an issue. Microsoft hasn't published exact limits, but users should be aware that the agent isn't a full ETL tool – mega datasets might still require traditional analysis tools or a database approach.

- **Supported Data Types:** Analyst supports many common file types (Excel, CSV, TSV, JSON, Power BI data via export, etc.)[\[7\]](#), but it may not support every possible format. Unusual proprietary data formats would need conversion before the agent can use them. Also, if a file is badly corrupted or password protected, Analyst can't handle that.
- **Reliability and Accuracy:** Despite the impressive reasoning, these agents are not infallible. They still **can make mistakes** or hallucinate information, especially if the source data has errors or if the AI extrapolates beyond the evidence. Microsoft itself cautions that while Researcher is more accurate than previous Copilot answers because it double-checks facts, it *"still possesses the ability, just like humans, to make mistakes"*. They stress that **verifying the agent's output is crucial for mission-critical decisions**[\[4\]](#)[\[4\]](#). In other words, users should not blindly accept everything – use the citations, check the logic (especially any analytical conclusions) and treat it as you would a report from a human junior analyst: mostly right but subject to review. The inclusion of source citations and code transparency is precisely to help with this verification.
- **Privacy and Data Compliance:** While the Copilot system is designed to respect privacy (data stays within your tenant and is not used to train the foundation model, per Microsoft's statements), organizations might have policies about how internal data is aggregated. For example, if some content is very sensitive, an admin might not want an AI agent pulling it into a summary that gets widely shared. Microsoft's Copilot Control System does allow IT admins to set which data connectors are allowed. So a limitation to note is that **admins might disable certain sources**. If Researcher isn't enabled on, say, your company's Confluence or CRM connector, it won't include that data. End users might sometimes find the agent doesn't pull something they expected because of such admin settings. In the admin center, there are toggles for enabling/disabling these agents for users and controlling their access to third-party connectors[\[10\]](#)[\[10\]](#). By default, if an organization has Copilot, these agents are **enabled** and available for users to discover[\[10\]](#), but that's subject to an admin's governance.
- **Throttling Under Load:** Microsoft has mentioned that users *"may experience throttling during heavy usage periods"*[\[10\]](#). This means if many people in your org (or across the service globally) are hitting Researcher/Analyst simultaneously, the service might make you wait or reject some requests until load decreases. Practically, this could show up as delays or being temporarily unable to use the agent. This is a common limitation for cloud services, but notable here because these agents presumably use a lot of compute per query.

- **Scope – Not a Real Human:** There are tasks that still require human judgment or creativity beyond what these agents can do. For instance, Researcher can compile information and even provide recommendations, but it won't *decide* on business strategy for you or take into account unquantifiable nuances (just as an AI, it lacks real-world experience or accountability). Analyst can analyze data given to it, but it might miss context that isn't in the data or known to the model. Also, if data is very outdated or incomplete, the agent can't magically fix that – it works with what it's given or can retrieve. It's important for users to understand the agents are **assistive tools** and not autonomous decision-makers. They augment human work; they don't replace the need for human review.

In summary, organizations deploying these agents should set the right expectations: they can substantially boost productivity and insight generation, but they have quotas, may require some patience on big tasks, and should be used with a verification mindset. Microsoft is actively improving them – features like broader language support for Analyst, higher reliability, and perhaps increased limits may come over time. But as of now, these constraints ensure users approach the agents' outputs thoughtfully.



User Feedback and Reception

The introduction of Researcher and Analyst has been met with a great deal of interest from early users. Feedback so far shows **significant positive impact on productivity**, coupled with some cautionary tales about the technology's current maturity.

On the positive side, **pilot users have reported dramatic time savings and improved outcomes**. In Microsoft's internal trials, users who had access to the Researcher agent found that tasks which used to take days of manual work could be completed in mere minutes^[5]. Product managers and

account managers in the pilot estimated that Researcher saved them on the order of **6–8 hours per week** each – essentially reclaiming almost a full workday that was previously spent on gathering and synthesizing information[5]. One product manager noted, *“It even found data in an archive I wouldn’t have checked. Knowing the AI searched everywhere—my meeting transcripts, shared files, the web—makes me trust the final recommendation much more.”*[5]. This quote highlights a couple of important perceived benefits: the agent’s thoroughness (searching “everywhere” one has access) and the trust that comes from seeing the comprehensive scope of its research. Users have described almost *“magical” moments* when the AI connected dots or uncovered useful info that they might have missed themselves[5].

Similarly, early business users have praised the **Analyst agent** for empowering them to get insights without waiting on data specialists. For example, a beta user in sales operations might report that they could quickly identify sales trends across regions using Analyst, whereas before they would have to ask an analyst for a report and wait a few days. The ease of just asking a question and getting charts and answers in return is frequently highlighted in demos and user discussions. The fact that Analyst shows its work (the code and calculations) has been cited as a confidence booster – it’s not a mystery how it got an answer, so users are more willing to trust and use those answers[2][1]. There’s a sense that these agents lower the barrier to data-driven decision making for non-experts, which is very welcome in roles like marketing, HR, or small business owners who can’t hire dedicated analysts.

External tech reviewers and community MVPs who tested the agents have largely echoed these positive sentiments. One Microsoft MVP, after hands-on testing, wrote that the Researcher agent *“blew my mind”* in how it *“transforms info gathering”* by synthesizing data from SharePoint, OneDrive, Teams, and the web using GPT-4 technology[6][6]. The ability to integrate third-party data (like Salesforce) into that flow was seen as a game-changer for getting a complete picture in research[6][6]. Another community blogger noted these agents feel *“like having a dedicated employee at your side ready to go, 24-7,”* aligning with Microsoft’s marketing that it’s akin to a permanent expert on call[3][3]. This kind of feedback underscores the value people see: it’s as if everyone suddenly has a research assistant or data analyst in their toolbox.

However, not all feedback is glowing—some early users have encountered **challenges and limitations in practice**, providing a reality check. An instructive example comes from a user who tried to use the Researcher agent in a complex real-world scenario (involving summarizing projects across a university with multiple specific requirements). This user wrote about how the **promise and hype met reality**: their query was quite detailed and they even supplied a Power BI

dashboard for data, but after half an hour the agent failed with a generic error[\[12\]\[12\]](#). This was understandably frustrating for them, as they expected at least partial results. The user identified possible causes – perhaps the request was too multi-faceted, combining very broad needs (strategic overview of many projects) with very granular demands (like SDG alignment for each project). The failure highlighted that the current agents can struggle with very **broad, multi-layered tasks** that a human might break into multiple reports or sessions. The user’s feedback pointed out the need for more **graceful handling**: ideally the AI should have either broken the task down or delivered whatever it could rather than nothing[\[12\]](#).

That said, it’s worth noting that Microsoft actively solicits feedback like this (the error message the user got even encouraged sharing feedback)[\[12\]](#). It shows the company is aware that these agents are a work in progress and is looking for edge cases to improve them. In forums and community discussions, Microsoft engineers have often responded that they’re working on improving reliability and scope coverage.

Another theme in user feedback is caution about **accuracy**. Even though Researcher cites sources, users have noticed it can occasionally cite something that, when clicked, doesn’t fully support the statement made (meaning the model may have slightly over-interpreted the source). This is a form of hallucination, albeit mitigated by the fact you can double-check it. Users are learning to treat the output as a first draft or an assistant’s work product – something to be reviewed. The good news is that many have found the accuracy to be much higher than generic GPT-based answers. Microsoft’s internal tests showed a nearly 70-90% improvement in various quality metrics over the base Copilot chat[\[5\]](#), thanks to the reasoning approach. But still, the advice “**trust, but verify**” has come up in user conversations, especially for critical analyses (no one wants to present a report to a client that has an erroneous data point).

General availability and enterprise adoption: As of June 2025 when Researcher and Analyst moved out of Frontier preview to general availability, all organizations with the Microsoft 365 Copilot license have access. Early adopters from the Frontier program (April-May) reportedly found these agents “highly effective for complex analytical work” and gave positive feedback that helped shape the final release[\[4\]](#). Customers in sectors like consulting, finance, and tech have been among the first to try these agents, given their natural fit in those knowledge-driven fields. The reception in press and media has been mostly positive, highlighting that Microsoft is pushing the envelope by integrating such advanced AI directly into mainstream productivity software, which could potentially change how people work on research and analytics tasks.

In summary, **user feedback confirms** that Researcher and Analyst agents can drastically reduce effort and uncover insights, essentially leveling-up what an individual can do without a support team. Users appreciate the transparency (sources, code) and see these agents as trustworthy aides, though not perfect. Some early hiccups in handling overly complex queries have been noted, and users recognize that these tools are still “1.0” and will evolve. The overall sentiment for many: a mix of *excitement* at the possibilities and *caution* to validate important outputs. As more users get access worldwide, we’ll likely see more creative uses and also more edge cases, but the foundational feedback is that these agents are indeed delivering significant value.

Market Positioning, Packaging, and Availability

Microsoft has positioned Researcher and Analyst as key differentiators in its business productivity suite, aiming to reinforce Microsoft 365 as *the* platform for integrated AI assistance at work. These agents are part of Microsoft’s broader **Copilot strategy** – which Microsoft describes as bringing “AI assistance into every layer of business operations”[\[2\]](#). Let’s break down how they are packaged, who can access them, their rollout timeline, and Microsoft’s messaging around them in the market.

Packaging and Licensing: Both Researcher and Analyst are not separate purchasable products; they are **included features of Microsoft 365 Copilot**. Microsoft 365 Copilot itself is an add-on license (at approximately \ \$30 USD per user per month) on top of Microsoft 365 (Office) subscriptions for enterprise customers[\[4\]](#). Therefore, to use these new agents, an organization must have opted into the M365 Copilot service. There is no direct consumer offering – these are aimed at enterprise and business users as part of the productivity cloud. In Microsoft’s product lineup, Copilot (with its various modes and now agents) is analogous to how Office was the essential toolkit of past decades; Microsoft is trying to make Copilot and its AI agents the new essential toolkit for modern knowledge work.

Within the Copilot experience, Microsoft has created an “**Agent store**” or **catalog** interface where pre-built agents like Researcher and Analyst can be discovered and enabled[\[10\]\[10\]](#). By default, every user with Copilot in an organization can find these agents listed under a “Built by Microsoft” category and start using them (unless an admin has restricted access)[\[10\]](#). This approach signals how Microsoft might introduce future agents (their own or third-party) – akin to an app store for

AI capabilities that admins can govern. Indeed, along with Researcher and Analyst, Microsoft has been previewing **Copilot Studio**, a toolset for organizations to build custom agents or “agentic” workflows themselves[\[1\]](#). This suggests Microsoft is positioning its platform to host not just these two, but many specialized agents (either built by Microsoft for general use, or by companies for internal use). Researcher and Analyst are essentially the flagship examples of what this new “**Frontier**” class of Copilot experiences can do, showcasing advanced reasoning in action.

Rollout Timeline: Microsoft announced Researcher and Analyst at its Enterprise AI event in late March 2025 (on Mar 25th, 2025)[\[1\]](#). They were introduced as “*the first two reasoning agents for Microsoft 365 Copilot*”, emphasizing their pioneering nature. Starting in April 2025, Microsoft made these agents available in preview to customers in the **Copilot Frontier program**[\[1\]](#)[\[10\]](#). Frontier is essentially an early access program for Copilot features where participating organizations (generally those who have already committed to Copilot licenses) could opt in to try out new features under development and provide feedback[\[10\]](#). Both Researcher and Analyst were labeled as “(Frontier)” in the Copilot interface during this phase to indicate their preview status[\[10\]](#). The preview rolled out gradually in April and May.

By June 2025, Microsoft announced the **general availability (GA)** of Researcher and Analyst for all Copilot customers[\[4\]](#). That means any user with a Copilot license (which at this time is primarily commercial Microsoft 365 E3/E5 customers who purchased the add-on) could now see and use these agents without special enrollment. The GA announcement also coincided with updates like the introduction of the 25-per-month usage cap, and the expansion of language support beyond just English[\[4\]](#). Essentially, June marked the transition from “preview novelty” to mainstream feature. Microsoft’s head of “AI for Work” Jared Spataro and even CEO Satya Nadella publicly touted the release with statements framing these agents as a new era of AI at work[\[2\]](#).



Mar 25, 2025: Agents Announced

Microsoft publicly unveils the Researcher and Analyst AI agents as part of Microsoft 365 Copilot, describing them as first-of-their-kind reasoning assistants for work.



Apr–May 2025: Frontier Preview

Early access rollout begins. Copilot customers in the Frontier program get to try Researcher and Analyst in preview (English-only, unlimited use during testing) and provide feedback.



June 2025: General Availability

Microsoft announces GA of Researcher and Analyst for all Microsoft 365 Copilot subscribers. Agents available to all licensed users (with usage limits and multi-language support introduced).



Beyond 2025: Expansion

Continued improvements expected – e.g. broader language support for Analyst, more specialized agents via Copilot Studio, and integration into more workflows – as Microsoft invests in “AI agents” as a key part of its productivity suite.

Geographic Availability: Microsoft 365 Copilot (and by extension these agents) is a **global service**, with Microsoft deploying it in datacenters that serve North America, Europe, and other regions. By GA, with support for 37 languages in Researcher, it's clear Microsoft aimed to make it usable by organizations in various countries across Europe, Asia, the Middle East, and the Americas^[4]. That said, initial preview customers skewed towards North America and Western Europe (where early Microsoft 365 Copilot adoption was highest). Markets like the Middle East would gain access as Copilot itself became available there – for example, support for Arabic would be key for adoption in that region. Given that Arabic is a widely used language, it's likely included in the 37 languages supported by Researcher at GA (though we'd have to confirm from Microsoft's documentation which languages exactly). Analyst's 8 languages probably included the biggest business languages (English, French, German, Japanese, perhaps Chinese, Spanish, Italian, Portuguese as a guess). Microsoft typically expands language support over time, so we can expect parity eventually.

One possible exception in availability is sovereign or government cloud environments. Sometimes new features reach Microsoft's commercial cloud first and only later the government or sovereign

clouds (due to extra compliance verifications). So large government or regulated industry tenants may not have had immediate access in some countries. But for the majority of commercial customers globally, by mid-2025 these agents are available.

Positioning in the Market: Microsoft's marketing and positioning of these agents emphasize them as a **competitive leap in workplace AI** that competitors are not yet offering within productivity suites. The messaging has been that *2025 is the "year of the AI agent"* and that Microsoft is leading the charge to bring agentic AI to everyday work[4]. CEO Satya Nadella tweeted that having Researcher and Analyst is like having *"a highly skilled expert on call for you 24/7"* inside your apps[2], reinforcing the idea that these aren't just chatbots, but expert collaborators. Microsoft is effectively positioning Copilot (with these agents) as **an AI that can do the higher-level thinking and analysis, not just clerical tasks**. This helps differentiate it from the first wave of AI assistants (like the initial Copilot which could draft emails or summarize documents but not perform deep reasoning).

In terms of market comparison, Microsoft is highlighting the **integrated nature** of their solution. In a press piece, they pointed out that these tools work "directly inside familiar tools like Word, Excel, and Teams"[2] – in contrast to some competitors where AI might be a separate chat interface or require switching context. They also stress **security and compliance** ("enterprise-grade security, privacy, and compliance" is a tagline in their materials), implying that using these agents within Microsoft 365 is safer for corporate data than using third-party AI tools externally[2]. The Copilot Control System and the fact that everything is within the Microsoft cloud is a selling point to CIOs.

Microsoft also launched a program called **Copilot Frontier** (as discussed) not just for early access but as a way to signal that customers can be part of co-developing these advanced AI solutions by providing feedback. This engages enterprise customers and makes them feel they have a hand in shaping the product – which is a smart move to increase adoption and loyalty.

Competitive Landscape: We'll dive deeper into direct comparisons in the next section, but from a positioning standpoint, Microsoft clearly sees these agents as a way to **maintain leadership against competitors like Google Workspace (with Duet AI) and emerging OpenAI offerings**.

Microsoft's advantage, which they underscore, is that they have an entire productivity ecosystem (Office, Teams, SharePoint, etc.) that the AI is natively plugged into[2]. This means their AI can take actions (like edit an Excel file, schedule a meeting, fetch an email) and use context that standalone AI services cannot easily get. Jared Spataro described these agents as delivering "advanced reasoning" in a way that *"Microsoft 365 is no longer just a helpful assistant... it can act as an*

on-demand subject matter expert"[3]. This positions Microsoft 365 Copilot not as a minor add-on, but as a transformational upgrade to Office that competitors would need to match.

Availability to End Users: For an end user at a company with Copilot, using Researcher or Analyst simply requires opening the Copilot panel (for example, in Teams or Office.com) and selecting the agent. If they don't see it, it could be the admin hasn't enabled the Agents or the user isn't licensed. Admins manage access via the Microsoft 365 admin center under Integrated Apps settings for Copilot agents[10]. By default, if you had Copilot enabled, the agents were available to install for users (they show up and the user can click to start using). In GA, Microsoft likely auto-installed or pinned them for users, as they referred to "pre-pinned agents" with usage limits[4] – suggesting these agents might appear by default in the Copilot experience.

Lastly, Microsoft's packaging includes these agents under the umbrella of "Wave 2" Copilot features (a term some blogs used), meaning after the initial rollout of Copilot's basic features in late 2023, these agents and some other enhancements (like image generation in Word, or memory and personalization features) are part of the next wave of improvements to keep adding value to the Copilot subscription[3]. This continuous enhancement strategy is how Microsoft is justifying the additional cost – by showing that Copilot is getting more and more capable over time (Researcher and Analyst being prime examples of new capabilities that make that \$30/month more worthwhile).

In summary, **Microsoft's market strategy** with Researcher and Analyst is to solidify Microsoft 365's position as the premier productivity suite with deeply integrated AI. They bundle these agents into the existing Copilot offering, leveraging their vast Office user base. The availability ramp went from a controlled preview to full release in just a couple of months, indicating Microsoft's confidence and urgency in getting these out (likely to stay ahead of rivals).

Geographically and linguistically, they pushed a broad support to capture global enterprises. From a messaging standpoint, Microsoft is saying: *we have not just AI, but AI that reasons* – effectively trying to leapfrog simpler AI assistants and anchor customers in their ecosystem with unique AI capabilities.

Comparison with Other AI Solutions

Microsoft's introduction of Researcher and Analyst agents comes amid a broader industry trend toward more autonomous, multi-step AI assistants (often referred to as "AI agents"). It's useful to compare these Microsoft agents to similar offerings or approaches from other vendors, such as

OpenAI's ChatGPT (and its tools), Google's Duet AI in Workspace, and others like IBM or Salesforce. We'll see where Microsoft stands and how these agents relate to or differ from the competition.

OpenAI (ChatGPT & Agents): OpenAI's ChatGPT is a well-known general AI assistant, and with the release of features like **Advanced Data Analysis** (formerly called Code Interpreter) and **Browse with Bing**, ChatGPT has gained capabilities somewhat analogous to Analyst and Researcher respectively. For instance, ChatGPT's Advanced Data Analysis plugin allows it to analyze files, run Python code in a sandbox, and generate charts – very similar to what Microsoft's Analyst agent does[\[13\]](#)[\[13\]](#). Likewise, ChatGPT can browse the web and summarize information (with the browsing feature), somewhat like a simpler version of Researcher. However, there are important differences. ChatGPT does these things in a more *ad hoc* way – you have to manually enable plugins, upload files, and ask it to do tasks, usually one at a time. It doesn't automatically integrate with your personal or enterprise files (unless you manually provide them), whereas Microsoft's agents natively hook into your Office 365 data (with permissions) without you having to upload anything. Also, Microsoft's Researcher is specialized to iteratively search *enterprise content* in addition to the web, which ChatGPT won't do out-of-the-box for a company's internal SharePoint, etc. OpenAI has hinted at "agentic" capabilities, and indeed in late 2024 there were reports that OpenAI was planning a new line of more autonomous AI agents for businesses, potentially at high price points (\$2k–\$20k per month for advanced research or coding agents)[\[2\]](#). These would aim to perform multi-step tasks. However, at the time Microsoft launched Researcher and Analyst, OpenAI's generative AI for end users was still largely ChatGPT and its plugin ecosystem, which requires more *user micromanagement* (the user often must prompt each step) compared to Microsoft's solutions that internally orchestrate multiple steps.

Another difference is **context and control**: Microsoft's implementation benefits from the Copilot control layer (ensuring enterprise security) and presumably uses the combination of OpenAI models plus Microsoft's own technology (like the Bing search index, Microsoft Graph data, and the new "Mu" model for on-device tasks in Windows, etc., though Mu is more for Windows as an aside). In essence, Microsoft has leveraged OpenAI's model strengths but built an integrated product around them tailored for enterprise workflows. OpenAI's ChatGPT, while powerful, is a general tool and not embedded in user workflows beyond what the user does manually (though OpenAI's ChatGPT Enterprise launched in 2023 allowed companies to connect internal data via tools like retrieval plugins, it's not as seamless as what Microsoft has done in M365).

However, it's worth noting that both Microsoft's agents and ChatGPT share a lineage in the sense that both owe their reasoning ability to OpenAI's model innovations. The difference is in integration: Microsoft's advantage is in leveraging the Microsoft 365 platform. OpenAI's advantage is being model-agnostic and directly fine-tuning for agent behaviors – indeed, the very “O3” models that power these Microsoft agents are OpenAI's creations^[5]. So one could say Microsoft's Researcher and Analyst are actually a showcase of OpenAI's latest tech (like GPT-4 and its successors) being applied in a specific domain. If OpenAI comes out with their own “Business Analyst GPT” or “Research GPT” as separate products, they'd likely operate outside of Office, which could be less convenient for Microsoft shops but potentially appealing to those who are not on Microsoft's stack.

Google Workspace Duet AI: Google has been integrating generative AI into its Workspace apps under the umbrella “Duet AI.” For example, Duet AI in Google Docs can help write content, in Gmail can draft emails, and in Google Sheets it can help generate formulas or summarize data. Google Sheets with Duet AI, around 2023-2024, introduced features like letting users ask questions about data (somewhat like a simpler Analyst), and autocompleting things like “generate a summary of this data” or creating charts based on prompts^[13]. However, Google's approach, at least initially, was more focused on assisting with single tasks in each app (like writing a paragraph, or filling a column with a formula) rather than orchestrating a multi-step reasoning chain that spans across data sources. For instance, Duet AI could help you analyze data in a Google Sheet you have open by generating an insight or pivot table, but it wouldn't automatically search across your Google Drive for relevant files, plus search the web, and then produce a multi-source research report – you would have to manually do parts of that.

Google has announced **Project Tailwind** and later its evolution in Google's NotebookLM, which is an AI that can ingest a set of documents and allow a research-oriented dialogue. That is conceptually similar to parts of what Researcher does (reading multiple user-provided docs to answer questions). But NotebookLM is still experimental and separate from Google Workspace's core apps. Meanwhile, Microsoft put Researcher directly into the Office workflow.

That said, **Google is not far behind in intent.** In 2023-2024, Google touted that its next-gen foundation model (Gemini) will enable more powerful capabilities in Workspace, possibly including more agent-like behavior. Google could integrate its AI with search and user data in similar ways given Google's expertise in search and cloud data. But as of mid-2025, Google's Duet

AI for Workspace was primarily offering generative writing help, automated slide generation, and chat assistance within apps, rather than an autonomous research or coding agent. Microsoft seized the narrative by giving these capabilities distinct names (“Researcher” and “Analyst”) and framing them as pushing boundaries.

Other Tech Companies: Several other players are in the “enterprise AI assistant” space:

- **Salesforce:** Has **Einstein GPT** and the Salesforce **Copilot** concepts for CRM, and they have capabilities to summarize Salesforce data, answer questions about pipeline, etc. That’s somewhat analogous to an Analyst agent but very domain-specific (sales data and CRM related tasks). Salesforce is also introducing an “Einstein Copilot Studio” for custom AI apps, similar to Microsoft’s Copilot Studio, which could allow building agents. But Salesforce’s focus is CRM and customer 360 data, not general productivity.
- **IBM and Watsonx:** IBM’s Watson Assistant and the newer Watsonx platform allow building conversational agents for business, and IBM no doubt is exploring agents that use retrieval and action. However, IBM’s offerings usually require more custom building by the client; they’re not pre-packaged for end-users like Microsoft’s. IBM might claim more customizability or data privacy by running models on your cloud, but Microsoft’s advantage is out-of-the-box utility.
- **AWS:** Amazon has Bedrock for foundation models and has integrated some generative AI in offerings like QuickSight (business intelligence) for natural language questions on data, which is a bit like Analyst but in a BI tool context. They also have CodeWhisperer for coding, not directly relevant here. Amazon doesn’t have an Office suite, so they aren’t a direct competitor in this specific use case.
- **Specialized tools:** There are startups and tools, like **Perplexity AI** or **Bing Chat (Enterprise mode)**, that provide research assistant-like capabilities. Perplexity, for example, does web search and text synthesis with sources, performing some multi-step reasoning when asked complex questions. It’s similar to **Bing Chat in Enterprise** which Microsoft offers (that uses GPT-4 with web and internal data if set up). In fact, one could consider that Microsoft had a precursor to Researcher in **Bing Chat Enterprise**, which allows employees to ask work questions and get answers with web sources but with privacy. However, Bing Chat Enterprise doesn’t directly dive into your emails or files – it only uses what you explicitly give it or what’s on the web. Researcher goes deeper by default into your internal corpus given permission. So Microsoft’s own lineup had different flavors: Bing Chat Enterprise (ad-hoc Q&A with web + user-provided text) vs. Researcher (multi-hop deep research with integration to M365).
- **Open-source agent frameworks:** In the broader AI community, things like **AutoGPT** and **LangChain** agents became popular for tech-savvy users to orchestrate LLMs towards goals. Those demonstrated what agentic AI could do, but required technical setup and

were not integrated with enterprise data out-of-box. Microsoft's approach is to provide that power without expecting customers to cobble things together.

Microsoft's Differentiation: The key differentiators for Researcher and Analyst as of 2025 are:

- **Deep integration with enterprise data and Microsoft 365 apps:** No competitor currently offers an AI agent that can seamlessly read your organizational emails, files, chats and also web info and then output a Word document or analysis that you can continue working on in Office. Google's closest would be if they deeply tie Google Drive/Workspace data to a future Gemini-powered agent, but that's in progress. Microsoft has a lead here due to Graph and Office integration. As one analysis noted, Microsoft's strength is that these AI tools *"don't sit off to the side—they're already living inside the apps people use daily."*[\[2\]](#).
- **Reasoning complexity:** Researcher and Analyst use iterative reasoning and even self-coded execution. Most other readily available AI assistants (ChatGPT included, by default) answer in one shot. Getting ChatGPT to do multi-step reasoning usually requires the user to prompt it step by step. Microsoft's agents automate that process internally. Google's Duet doesn't (yet) write and execute code by itself, for example, whereas Analyst does. This puts Microsoft at the frontier of what is sometimes termed "Autonomous AI agents" in a practical user-facing way.
- **Responsible AI and compliance:** Enterprises might prefer Microsoft's offering because Microsoft can assert control and monitoring (via the Copilot control system, logging of agent actions, adherence to data compliance). Using something like ChatGPT directly might raise compliance questions for sensitive data. Microsoft emphasized that everything these agents do is audited and controllable by IT[\[1\]](#), and that agents follow the same permission model as the user (they can't access what the user can't). This enterprise-ready framing is a comfort that not all competitors have out of the box.
- **Ecosystem and extensibility:** With Copilot Studio, Microsoft signals that companies will be able to build their own specialized agents or extend these. For instance, a company could create a "Legal Researcher" agent that only looks at internal legal docs and certain databases. While others like OpenAI also allow building custom GPTs or agents, Microsoft gives a more turnkey solution for that integrated into the same Copilot environment.

It's clear Microsoft also views these agents as a way to **keep Microsoft 365 valuable** in an AI-first world. If AI assistants that sit outside (like say, a hypothetical future "Google Assistant for Work" or a standalone AI from OpenAI) became powerful enough, one might question the value of traditional software suites. By embedding the best AI into M365, Microsoft is ensuring customers stick with or even expand their use of Microsoft's ecosystem rather than using external AI tools.

Competitive Response: We can anticipate competitors responding. Google will likely enhance Duet AI to do more multi-step, connected tasks – possibly integrating Google search and user data more deeply. For instance, one could imagine a future Google agent that, when asked a question,

automatically searches your Gmail, Google Drive, and the web and gives a comprehensive answer (which would parallel Researcher). Google certainly has the pieces (Search, Gmail, Drive, Analytics), so it's likely a matter of time. OpenAI might offer businesses the ability to plug ChatGPT into their data with more autonomy (they launched "ChatGPT Enterprise" and "ChatGPT with custom GPTs" which could be configured similarly, but it requires effort and isn't inside your normal tools).

There are also other big tech efforts: for example, Meta AI introduced some agents (but mostly consumer-focused for now), and smaller niche companies offering tools for literature review (like Semantic Scholar's AI, or Elicit.org for research). Microsoft's advantage stands that they have distribution – millions of Office users who can get these features without switching context.

To sum up, **Researcher and Analyst vs others:**

- They are ahead in **integration** (compared to ChatGPT which is outside your work environment).
- They are ahead in **specialization and chain-of-thought** (compared to current Google Workspace AI features).
- They face competition from **alternative approaches** (like if a company prefers to use ChatGPT with their own data, or if they use a different productivity suite).
- They prompt competitors to accelerate: indeed, the presence of these agents "helps Microsoft keep up against the competition, which is also aggressively pursuing agents" [\[4\]](#), and arguably in some areas vaults Microsoft slightly ahead for the moment.

In the bigger picture, this aligns with a trend identified in late 2024: tech leaders declared that 2025 would be the "year of **agentic AI**" where AI moves from just generating text to taking actions and doing multi-step tasks autonomously [\[4\]](#). Researcher and Analyst are concrete manifestations of that trend in a mainstream product. Competitors will certainly have their own takes (OpenAI's planned agents, Google's continuous Workspace AI improvements, etc.), but Microsoft has set a high bar by delivering this inside a widely-used platform early in 2025.

Conclusion

Microsoft's Researcher and Analyst agents exemplify the next generation of AI assistance in the workplace: **domain-specialized, reasoning-powered agents** that go beyond drafting emails or documents, to actually *analyzing data, synthesizing knowledge*, and providing expert-like support. Architecturally, they blend OpenAI's cutting-edge language models (the "O3" series, related to

GPT-4) with Microsoft's robust graph, search, and tool execution framework – resulting in a hybrid that can both **think** through problems step-by-step and **act** (by fetching information or running code).

We've seen that Researcher can serve as an ever-ready research analyst, capable of writing a detailed business report that pulls from your SharePoint, your emails, and the web, citing every quote. Analyst, on the other hand, can be your personal data scientist, able to crunch numbers from multiple Excel files, find patterns and even generate Python code to solve a data problem, then hand you the results in plain language with a chart. These are sophisticated capabilities that, until recently, were in the realm of skilled human specialists or complex manual processes.

The use cases are compelling: from preparing for key meetings with comprehensive briefs, to generating on-the-fly business insights without waiting weeks for an analytics team. Early user feedback indicates real productivity gains – hours saved and better decisions made – thanks to more thorough information gathering and analysis. Users also appreciate the design choices of Microsoft's agents: the transparency (seeing sources and code) which fosters trust, and the interactive clarification which leads to more relevant outcomes.

However, it's equally clear that we are in the early days of such AI agents. They come with limitations: query caps, occasional slow or incomplete responses on extremely complex asks, and the enduring need for human oversight to catch any mistakes. Microsoft acknowledges these limitations and seems committed to improving the agents' capabilities, using techniques like reinforcement learning to make them smarter and more reliable with time[\[5\]\[5\]](#). We can expect updates that will allow deeper integration (e.g., chaining multiple agents seamlessly, as hinted with specialized agents working together) and more fine-grained user or admin control over sources and behavior[\[5\]\[5\]](#).

In terms of market impact, Researcher and Analyst strengthen Microsoft's hand in the battle of workplace AI. Google and others will likely accelerate their own agent offerings, but Microsoft's strategy of embedding these into the daily tools of work (Word, Excel, Teams) and tying them to the user's actual data and context is a strong one that competitors will have to match. It also potentially drives more adoption of Microsoft 365 itself, as organizations see the value of an AI-enhanced workforce. Microsoft is effectively adding an AI "brain" on top of the Office suite that can leverage everything in an enterprise's digital estate – something very hard for a point-solution competitor to replicate without similar integration.

So, how do these agents relate to similar AI services? They are both a response to and a lead over the current landscape. They incorporate ideas seen in places like ChatGPT (for coding assistant) or

Bing (for web Q&A) but unify them in a purposeful way for the enterprise environment. Microsoft has set a precedent by giving these agents distinct roles and names, making it easy for users to understand what tool to use for what job (research vs data analysis), whereas other ecosystems might require more user-driven setup (e.g., picking the right plugin or prompt).

To sum up, **Researcher and Analyst represent a significant innovation in AI for productivity.**

They illustrate a future where every professional has on-demand access to what feels like a team of assistants – one that can scour all your information and knowledge to answer complex questions, and another that can crunch numbers and illuminate data-driven insights – all within seconds and all within your trusted work environment. This augments human workers' capabilities, allowing individuals and small teams to punch above their weight in terms of research and analytical output.

However, with great power comes the responsibility to use it wisely: users must learn to craft good questions, to review AI outputs critically, and to collaborate with these agents (providing clarifications or additional inputs) to get the best results. Organizations will also navigate how to govern these tools (ensuring they align with compliance and that employees understand their best use).

In conclusion, Microsoft's Researcher and Analyst agents are trailblazing examples of AI's evolution from simple assistants to **intelligent collaborators**. They are built on a fusion of advanced AI modeling and practical enterprise integration, use cases show they can revolutionize daily work from strategy formulation to data analysis, and initial reception is very promising. While not perfect yet, they are continually improving and mark a turning point in how AI can contribute to knowledge work. As competitors develop their own versions and Microsoft expands its roster of agents, one thing is clear: *the era of AI agents in the workplace has truly begun*, and Microsoft's Researcher and Analyst are at the forefront of this new era[4][4].

References

- [1] [Microsoft Announces Researcher and Analyst Agents for Microsoft 365 ...](#)
- [2] [Microsoft launches CoPilot AI Agents, 'Researcher' and 'Analyst,' aim ...](#)
- [3] [Inside Copilot's Researcher and Analyst Agents](#)
- [4] [Microsoft's new AI tools: What "Researcher" and "Analyst" mean ... - Neowin](#)
- [5] [Researcher agent in Microsoft 365 Copilot](#)
- [6] [I Tested Microsoft's AI Researcher Agent \(And It Blew My Mind\)](#)
- [7] [Analyst agent in Microsoft 365 Copilot](#)

- [8] [Get started with Analyst in Microsoft 365 Copilot](#)
- [9] [3 practical ways small businesses can use Researcher and Analyst agents](#)
- [10] [Microsoft 365 Copilot Frontier experiences becoming available](#)
- [11] [Get started with Researcher in Microsoft 365 Copilot](#)
- [12] [When AI Research Agents Hit Reality: My Microsoft Researcher Experience](#)
- [13] [Which AI Can Create Excel Spreadsheets? Here's What You Need to Know ...](#)