



Copilot Studio Agents Report Interpretation Guide

Last updated 7/29/25

The content in this interpretation guide is intended solely to help you better understand the structure and insights provided in the report. The charts here are based on demo data and your insights may look materially different.

Introduction

This interpretation guide is designed to help users better understand the results in their Copilot Studio Agents report. It highlights some of the most common trends and patterns observed across organizations using the Copilot Studio agents report.

Detailed Slide Interpretation

Walks through each section in the Copilot Studio Agents report in detail explaining the page's purpose, how to read it, and suggestions for further analysis. Refer to this section as you review your report to ensure you have a clear understanding of the data on each slide and to get ideas for next steps.

Questions?

Consult the FAQ (slide 3-5) and Glossary (slide 6 & 7) for additional helpful resources

FAQ

Q: What is the Copilot Studio Agents Report?

A: The Copilot Studio Agents Report is a Power BI-based analytics tool available in Viva Insights. It helps organizations understand the adoption, usage, and impact of agents built in Microsoft Copilot Studio. It provides insights into:

- Top agents by usage
- Key performance indicators (KPIs) like sessions, satisfaction scores, and success rates
- Agent-assisted hours and topic distribution
- Engagement trends across deployment channel

Q: What's included in this report?

A:

- Agents built in Microsoft Copilot Studio and deployed through the following channels: custom website, demo website, Microsoft Teams and Microsoft 365 Copilot, mobile apps, Facebook and Azure bot services channels such as Skype, Cortana, Slack, Telegram, Twilio, Direct Line Speech Email, and GroupMe
- Metrics on agent engagement and outcomes
- Breakdown of sessions by topic and user interaction patterns

Q: What's not included in this report?

A:

- Agents built using the Agent Builder in Microsoft 365 Copilot
- Agents built using Copilot for Microsoft 365 within Copilot Studio (Declarative Agent)
- Agents that use generative AI orchestration for responses (conversational or autonomous agents)
- Agents built outside of Copilot Studio

FAQ

Q: What are the prerequisites to use the report?

A: To access and run the report, your tenant must:

- Have at least 50 Copilot licenses
- Include at least one Microsoft Copilot Studio license (standalone, pay-as-you-go, or bundled)
- Have at least one published and usage of the agent in a "Production" environment
- Use Power BI Desktop (June 2022 or newer)

Q: Where to find the report?

A:

- Navigate to the Viva Insights portal where you can run advanced queries.
 - Leaders will be able to see " Copilot Studio Agents report " under reports tab – if agent data is available for 1 month
- Analyst will be direct to the "Advanced Analysis" page where can view the report with 1 click or select "Create Analysis" - This opens the Power BI query builder interface.

What information do you need to run this report?

A:

- Time Period: Defaults to Last 1 month. To start, you can only run an analysis covering the last month, but eventually you'll be able to select a wider time period as Microsoft Copilot Studio collects more data.
- Predefined template metrics: Preselected metrics will appear as gray tags. These metrics are required to set up the Power BI report and you can't remove them. You can't add other metrics for this report.
- Which agents you want to include in the query; you can add filters to narrow down the agents in scope for your report. You can filter by agent name, agent surface, and agent type.
- Select which agent attributes you want to include in the query: The attributes are selected automatically. Once the query runs, you can use these attributes to group and filter the reports.

FAQ

Q: How are Copilot Studio Agents report filter user types defined?

A: The report backend (Power BI or Viva Insights) utilizes telemetry from Copilot Studio to run the data science agent model, determining whether an agent is internal or external.

These classifications are surfaced in the report as filterable dimensions which distinguishes between how and by whom agents are used. Here's how these user types are structured and applied:

- **Internal Users:** These are employees within your organization. Internal usage typically includes agents deployed for:
 - IT support
 - HR assistance
 - Internal knowledge retrieval
 - Sales enablement for internal teams
 - Purpose: Improve employee productivity, streamline internal workflows, and reduce time spent on repetitive tasks.
- **External Users:** These are customers, partners, or users outside your organization. External usage includes agents embedded in:
 - Customer service portals
 - Public websites
 - Partner support channels
 - Purpose: Enhance customer experience, scale service delivery, and reduce support costs

Glossary

See here for a full description of all Insights metrics and terms:
<https://learn.microsoft.com/en-us/viva/insights/advanced/analyst/templates/copilot-studio-agents>


Active agents	Total number of agents with at least one session.
Agent assisted hours	Estimates the time employees save by using an agent to perform tasks. It is calculated by the number of agent sessions and a time savings multiplier, which can be customized in the Calculator.
Agent assisted value	An agent’s assisted value is calculated by looking at the agent-assisted hours and applying an average hourly rate, which can be customized in the Calculator.
Agent satisfaction score	At the conclusion of an agent session, users receive a survey asking them to rate their experience from 1-5. If users complete the survey more than once (due to multiple uses of the agent), only the most recent response per session counts. The average of all survey responses is reported as the overall satisfaction score.
Abandoned session	An interaction is considered "abandoned" when a user exits the session. This includes before their goal is achieved—typically after multiple failed attempts or confusion - or when a user got what they needed, but didn't give a thumbs up or say "thank you" or "goodbye.” For example, if an employee asks about a client's refund amount and does not respond after the agent replies with the refund status, this is considered abandonment.
Agent session outcome	Percentage of engaged agent sessions, grouped by outcome. An agent session can either be resolved, escalated, or abandoned.
Agent session resolution rate	Percentage of engaged sessions across all deployed agents that were resolved (i.e., neither abandoned nor escalated).
Category	This is a way to group agents according to their functional purpose or where they are deployed.
Conversation	An ongoing interaction between a user (or group of users) and an agent across a specific channel (e.g., Teams, web, mobile).
Engaged agent sessions	Total number of sessions with an agent during the selected time period. An engaged session means the user clearly intended to get something valuable from the conversation.
Escalated Session	An interaction is escalated when the agent cannot resolve the issue and hands off the session to a human agent or another system.

Glossary

See here for a full description of all Insights metrics and terms:
<https://learn.microsoft.com/en-us/viva/insights/advanced/analyst/templates/copilot-studio-agents>

Interaction	Can be a message, command, or input from the user that the agent responds to.
Knowledge Source References	<p>A knowledge source is any kind of data—structured or unstructured—that Copilot can use to answer questions. This includes things like: PDFs, Word docs, Excel files, Dataverse tables, SharePoint sites, Websites (internal or public), Custom APIs, and Microsoft Graph-connected enterprise data.</p> <p>Whether the session uses a knowledge source or not impacts the way the report calculates the assisted hours for the session:</p> <ul style="list-style-type: none">• With knowledge sources: Each source used saves about 6 minutes. If two sources are referenced, that’s 12 minutes saved.• Without knowledge sources: If the session is marked resolved, it counts as 6 minutes saved.
Number of sessions per topic	Total number of engaged sessions per topic over the selected time period.
Resolved session	A session is marked as "resolved" when the agent successfully identifies the user’s intent and completes the task or provides a satisfactory answer.
Session	A session begins when a user initiates a chat with a Copilot agent, or the agent starts the conversation. It ends when the interaction concludes, whether it's resolved, escalated, or abandoned.
Topics	Topics are defined during the agent setup process in Microsoft Copilot Studio. They are the foundational elements that determine how an agent understands and responds to user input. Topics can be thought of as the skills of an agent. These topics are like distinct conversation routes that, when combined, enable users to experience a dialogue that feels smooth and natural (e.g., an HR agent identifies the main topic of a user's question—such as expenses, payroll, attendance, or other areas).
Usage type	Classification for segmenting and analyzing agent activity. When the agent is mainly used by the public or external customers, we classify it as <i>external</i> . This includes agents embedded in customer service portals, public websites, and partner support channels. An agent is considered <i>internal</i> when it's mainly used by employees within an organization. These agents typically help out with tasks like IT support, HR policies, internal knowledge bases, and sales tools for internal teams.

Detailed Slide Interpretation



Agent overview

1

How many agents were active during this time period?

INTERPRETATION: An *active agent* is defined as any agent where a user has participated in at least one engaged session during this time period.

2

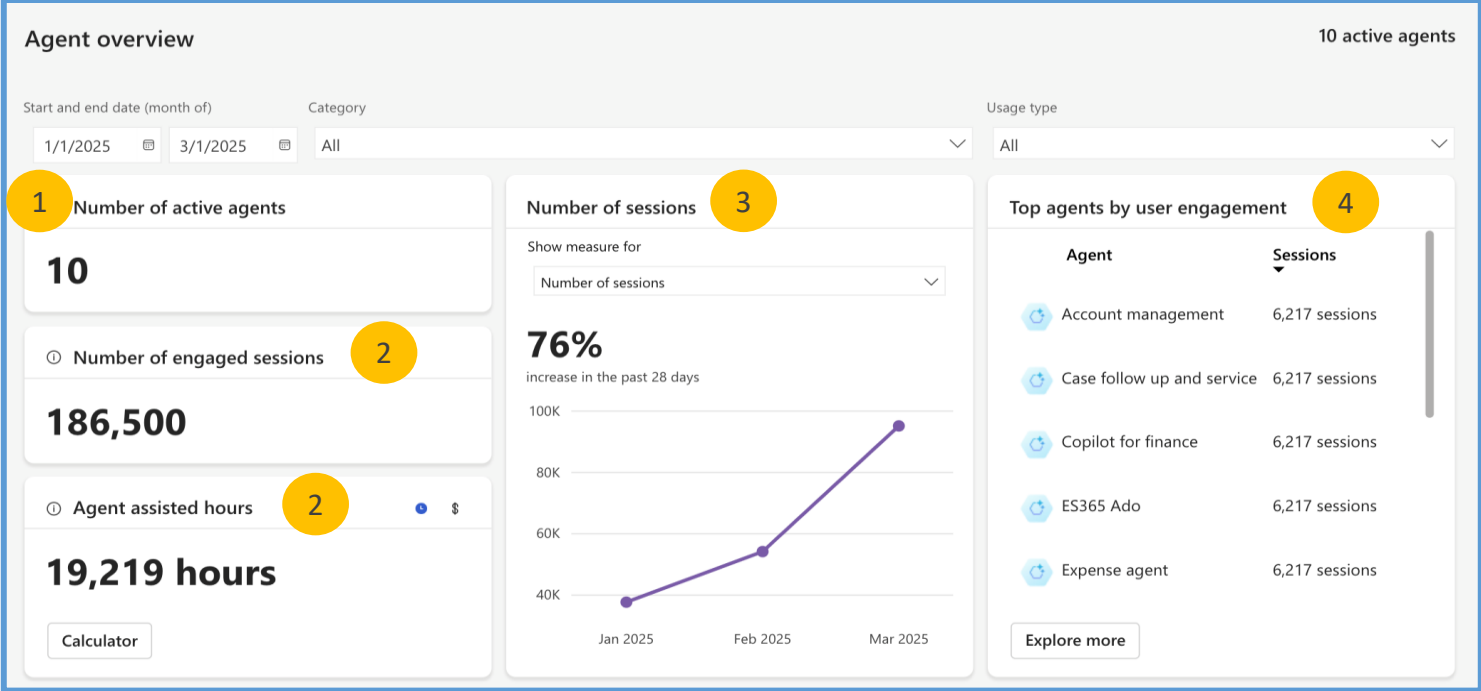
What is the total number of engaged agent sessions – and how much time or money have those sessions saved?

INTERPRETATION: The *total number of engaged agent sessions* indicates the number of times users engaged with the agents identified in #1.

To determine the assisted hours, the report looks at the type of request and makes an estimate. For example, requests with multiple knowledge source references would be estimated to have a higher time savings than a simple question and answer. By default, the report assumes 6 minutes per knowledge source reference OR session without a knowledge source reference; to view this breakdown or to adjust this assumption, click the “Calculator” button.

Alternatively, you may choose to view the impact by assisted value by clicking the dollar sign (\$) icon. The same calculation takes place as for assisted hours but instead uses an estimate of \$72 per hour of impact. As with assisted hours, this can be updated by clicking on the “Calculator” button.

GUIDANCE: Whether looking at assisted hours or assisted value, it is important to recognize that these are estimates to help you better understand the relative impact of agents through efficiency gains and time savings.



3

How have the number of active agents, number of sessions, and agent assisted hours have changed over time?

Use the dropdown to toggle between the metrics.

INTERPRETATION: The trend lines illustrate the change in key metrics over time.

GUIDANCE: We would like to see each of these trend lines growing as time progresses, which would indicate not just a growth in the number of agents, but in agent usage and impact as well.

RECOMMENDATION: Pay attention to the percentage numbers at the top of each chart as well, as this indicates the change over the previous month. Look at the percentages for number of sessions and assisted hours, and how they compare to the number of agents; if the number of agents is growing faster than these usage and impact metrics, it may be a sign to focus more on promoting and enabling the existing agents.

4

What are the top five agents based on monthly usage?

RECOMMENDATION: These are the most popular agents currently in use at your organization. Click on the “Explore more” button to view more in-depth metrics and view a full list of available agents.

Explore More: Top agents used across the organization

1

How should you use the number of engaged sessions, satisfaction score, and resolution rate together for Actionable Insights?

		Satisfaction Score	
		Low	High
Resolution Rate	Low	<p>This agent is not delivering value</p> <ul style="list-style-type: none">Review the agent's configuration to confirm it has been set up correctly.Prioritize for redesign or retirement.Interview users or review logs to understand failure points.Consider whether the agent's purpose is still relevant.	<p>Users like the experience but aren't getting what they need</p> <ul style="list-style-type: none">Investigate if the agent is too limited in scope or lacks access to needed data.Expand capabilities or improve escalation paths.Consider retraining the model or refining intents.
	High	<p>Tasks are being completed, but the experience may be frustrating or confusing</p> <ul style="list-style-type: none">Review the agent's configuration to confirm it has been set up correctly.Review conversation logs for friction points (e.g., tone, length, complexity).Improve User Experience elements like clarity, empathy, or speed.Consider user feedback mechanisms to pinpoint dissatisfaction.	<p>This agent is performing well and meeting user needs</p> <ul style="list-style-type: none">Use it as a benchmark—analyze what makes it successful (e.g., tone, flow, content).Consider scaling or promoting it more broadly.Replicate its design patterns in underperforming agents.

Start and end date (month of)		Category	Usage type	Agent
1/1/2025	3/1/2025	All	All	All
Average monthly metrics				
Agent	Engaged sessions	MoM % change in engaged sessions	Satisfaction score	Session resolution rate
Account management	6,217	+76%	4.3/5	87%
Case follow up and service	6,217	+76%	4.3/5	87%
Copilot for finance	6,217	+76%	4.3/5	87%
ES365 Ado	6,217	+76%	4.3/5	87%
*Accuracy and completeness of the session resolution rate are impacted by topic configurations, which can be updated by agent makers.				
Learn more				
Explore more				

GUIDANCE: Use the table to evaluate the effectiveness and satisfaction of your agents and determine potential action steps. By analyzing successful agents, you can use that insight to boost the performance of those that aren't doing as well.

Importantly, also consider the session volume in your evaluation; is the agent being used enough to draw meaningful conclusions? If you have agents with low usage, you might want to think about promoting them more, retiring them, or reworking their design.

RECOMMENDATION: You can filter the agents on this page using the Category, Agent name, and Usage type (i.e., internal or external to the organization).

Use the "Category" field (e.g., Sales, Support, Internal) to compare performance within similar agent types. This helps avoid unfair comparisons and surfaces best practices within each domain.

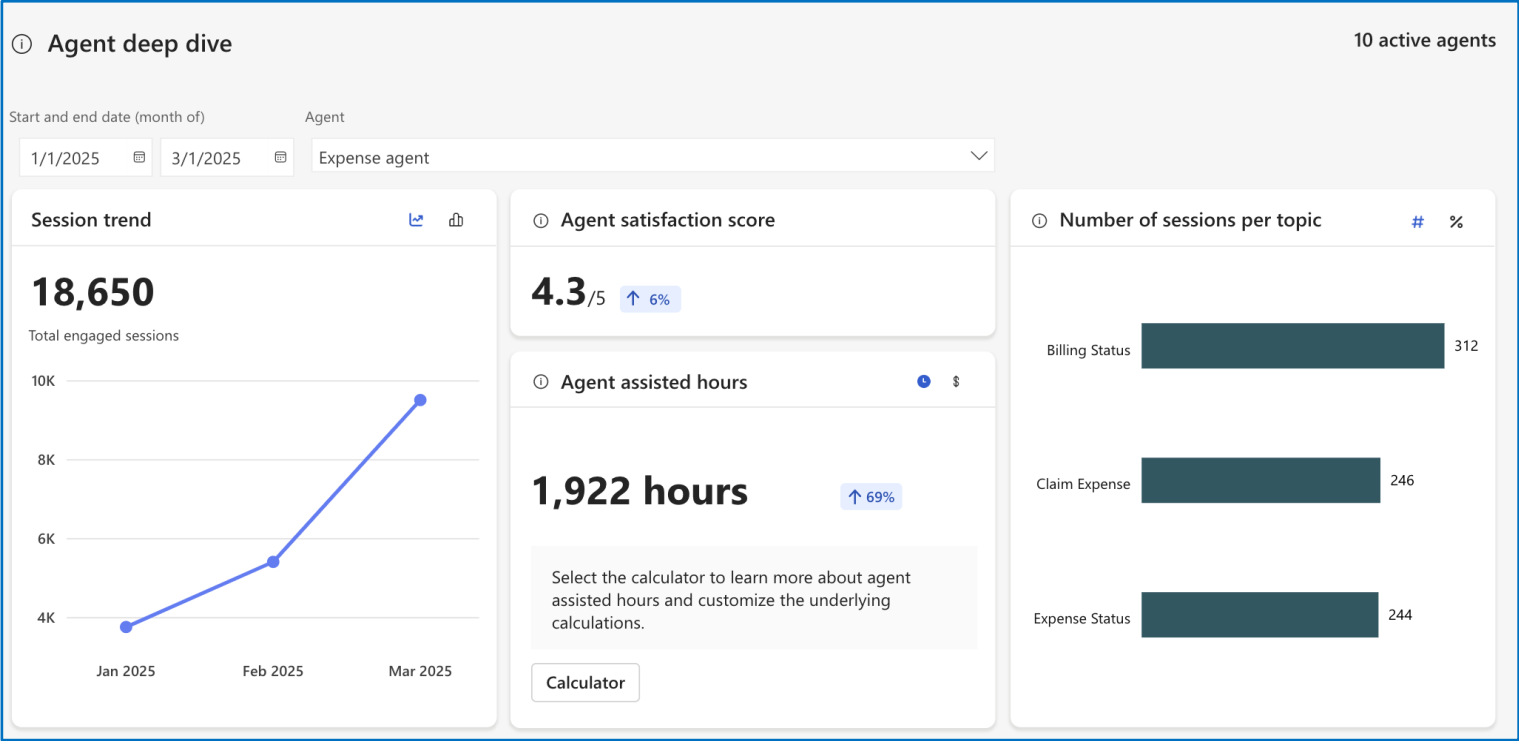
What can we learn about each individual agent’s performance?

When you navigate to the deep dive page in the report, you're entering a detailed analytics view that allows you to:

- Drill into specific KPIs for individual agents
- View engagement metrics like:
 - Sessions over time
 - Resolution rates
 - Abandonment and escalation rates
 - Satisfaction scores
 - Analyze agent assisted hours and topic distribution to understand adding value

GUIDANCE: Using the Agent dropdown, select an agent to explore further. Use the information on the following pages to learn more about:

- Which agents are performing well?
- Which ones may need optimization or retirement?
- How does the agent usage align with its intended goals?



Agent deep dive (2/5)

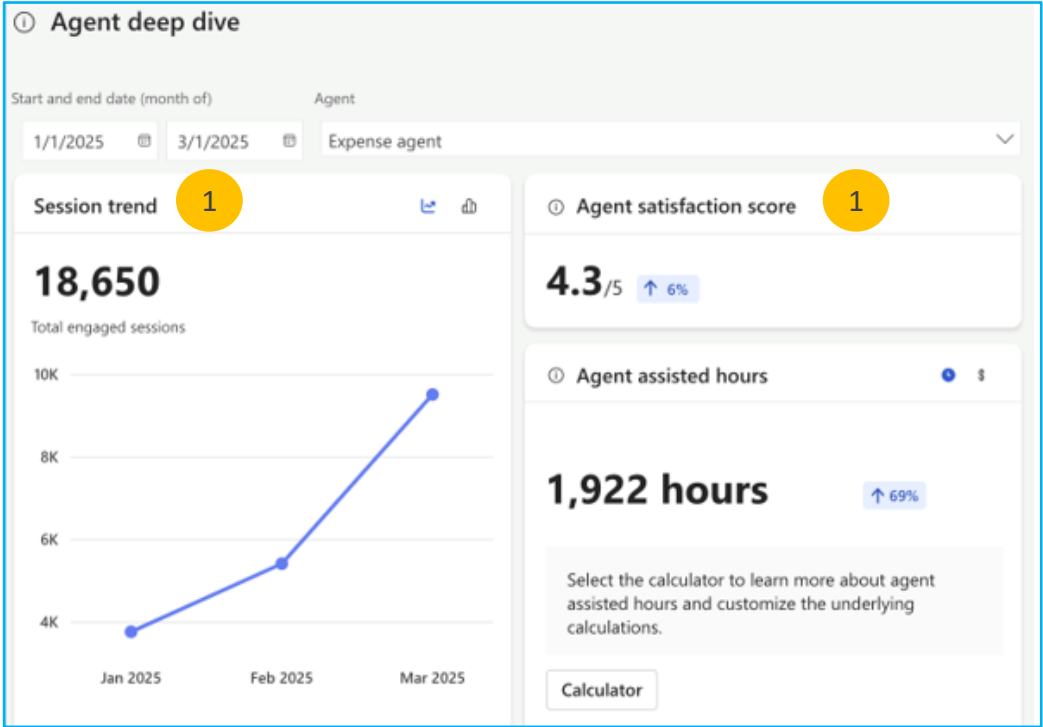
1

How have usage and satisfaction changed over time for this agent?

INTERPRETATION: These two metrics illustrate the number of people using the agent over the last three months, as well as the average agent satisfaction rating.

GUIDANCE: Use these sections to gain a high-level view on the success of the specific agent. Based on the trends you see, consult the table below for potential actions to take.

Trend	Potential Interpretation	Recommendation
Rising session volume	Increased adoption or successful rollout	<ul style="list-style-type: none">Investigate which agents or topics are driving growth.Reinforce successful patterns and replicate them across other agents.Review satisfaction scores to ensure quality is keeping pace with volume.
Declining session volume	Possible user disengagement, technical issues, or seasonal drop-off	<ul style="list-style-type: none">Check for recent changes in agent behavior, broken triggers, or user feedback.Re-engage users with targeted communications or improvements.Review satisfaction scores to identify if declining quality is contributing to the drop.
Flat or stagnant volume	Stable usage, but possibly underutilized potential	<ul style="list-style-type: none">Consider promotional efforts, training, or expanding the agent’s capabilities to dive more valueAnalyze satisfaction scores to determine if users are content or indifferent.
Low Volume	Very few sessions may indicate low adoption or discoverability.	<ul style="list-style-type: none">Promote the Agent in internal communications (e.g., newsletters, Yammer, Teams).Educate and train usersReview satisfaction scores from the limited sessions. Positive scores suggest the agent is on track and you may need to increase awareness. Negative scores indicate improving the agent before wider deployment.
Sudden spikes or dips in volume	Could indicate a campaign, outage, or a new feature rollout	<ul style="list-style-type: none">Coordinate with internal events or releases.Use these insights to plan future launches or mitigate issues.Check satisfaction scores to assess user sentiment during the spike or dip.


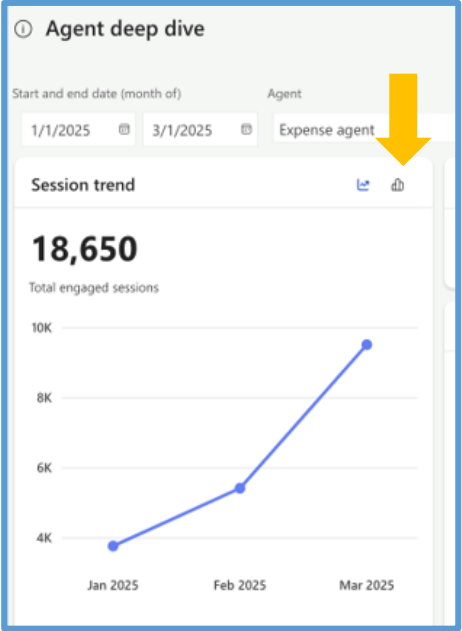


2 What are the Session outcome trends for agent maturity?

(Click on the bar graph in the upper right corner of the section trend to toggle to Session outcome trend)

GUIDANCE: It's great to see a positive trend where we're moving away from "Abandoned" towards "Resolved" over time. Ideally, satisfaction scores should increase as resolution rates go up and decrease when we see more cases of abandonment or escalation. Think of these scores as a way to confirm or question our volume-based interpretations.

Trend	Potential Interpretation	Recommendation
Increasing Resolved Sessions	Your agent is effectively handling user prompts/queries.	<ul style="list-style-type: none">Reinforce successful patterns—analyze which topics or flows are contributing to high resolution and replicate them.Review satisfaction scores to confirm that resolution quality aligns with user expectations.
Low Resolution Rate	A low percentage of successful outcomes suggests the agent isn't meeting user needs.	<ul style="list-style-type: none">Review unresolved sessions to identify common issues like misunderstood intents, missing content, or ineffective responses.Enhance training data and refine conversation flows to better support user goals.Analyze satisfaction scores for unresolved sessions to understand user sentiment—low scores may indicate frustration, while neutral or high scores could show users found value even without formal resolution. Use these insights to prioritize improvements.
Rising Escalations	The agent is hitting limitations in handling prompts/queries.	<ul style="list-style-type: none">Review escalation triggers. Are they firing too early?Could the agent be trained to handle more before escalating?Review the agent's configuration to confirm it has been set up correctly.
High or Growing Abandonment	Users are leaving before getting help—possibly due to confusion, delays, or irrelevant responses.	<ul style="list-style-type: none">Audit conversation paths. Look for friction points or dead ends.Consider configuration such as adding fallback logic.Review the agent's configuration to confirm it has been set up correctly.
Persistent Abandonment Outcomes	Sessions are starting but not engaging—users may be testing or unsure how to interact.	<ul style="list-style-type: none">Improve onboarding, configuration – fallback logic or use experience cues to guide users into meaningful engagement.Review the agent's configuration to confirm it has been set up correctly.



3

How much time is this agent saving across the organization?

INTERPRETATION: This number provides an estimate of the hours saved by this agent.

GUIDANCE: Click on Calculator to swap out the default 6-minute multiplier (based on industry research) with values from your own time trials and customer estimates. This way, you can tailor the metric to better reflect real-world impacts, ensuring a more accurate measurement.

Alternatively, if you prefer to view results as a dollar value, choose the \$ toggle to get an estimate of the dollars saved by this agent. As with assisted hours, you can customize the default value of \$72 per hour by clicking on the Calculator button and choosing “Value calculator.”

RECOMMENDATION: Identify the agents that contribute the most to time savings. Focus on optimizing or scaling those high impact agents. Link assisted hours to business outcomes like cost reduction, CSAT improvement, or support deflection.

Agent deep dive

Start and end date (month of)

1/1/2025

3/1/2025

Agent

Expense agent

Session trend

18,650

Total engaged sessions

10K

8K

6K

4K

Jan 2025

Feb 2025

Mar 2025

Agent satisfaction score

4.3/5

↑ 6%

Agent assisted hours

3

1,922 hours

↑ 69%

Select the calculator to learn more about agent assisted hours and customize the underlying calculations.

Calculator

Agent assisted hours metric definition

- ☒ Hours calculator
- ☐ Value calculator

This metric estimates the time employees save by using an agent to perform tasks. It factors in the number of knowledge source references, sessions without knowledge sources, and a time savings multiplier. This metric helps businesses understand efficiency gains and time savings from deploying agents.

Agent assisted hours calculator

- Minutes saved multiplier

6
- Average hourly rate

72

Measure	Knowledge source references	Sessions without knowledge sources	Minutes saved multiplier	Agent assisted hours
Agent assisted hours	10,211	9,008.0	6.0	1,922

(Knowledge source references + Sessions without knowledge sources) x Minutes saved multiplier = Agent assisted hours

About this estimation

What are people using this agent most often for?

INTERPRETATION: The variety of topics that come up during a single user session is a key measure of how complex the conversation is, how well the agent navigates it, and how satisfied the user feels about getting their task done.

GUIDANCE: Topics define the conversation paths within an agent. They represent specific intents or scenarios that the agent can handle, such as answering FAQs, completing tasks, or troubleshooting issues. Without well-structured topics, the agent cannot deliver meaningful or accurate responses.

High frequency across numerous topics could indicate a dynamic agent eager to explore various areas; this is beneficial for exploratory or research-focused agents meant to cover multiple topics. However, it might also signal issues like poor topic design, unclear guidance, or excessive fallback triggers.

RECOMMENDATION:

To analyze topic effectiveness, click on an individual topic under "Number of Sessions Per Topic" to explore the details further. Some suggestions for further analysis include:

1. **Track Resolution Rates by Topic:** Observe how frequently sessions on a specific topic conclude with a "Resolved" status, indicating the topic is well understood and effectively handled.
2. **Monitor Escalation and Abandonment:** Identify topics that lead to escalations or users giving up, signaling possible content gaps or user frustration. Compare these trends with satisfaction scores to determine if users are dissatisfied or simply need more guidance.
3. **Evaluate Engagement Metrics:** Assess the duration of user engagement with each topic and their interaction with multiple steps or follow-up prompts. High engagement and positive satisfaction scores typically signify a well-designed topic flow.
4. **Use Satisfaction Scores for Sentiment Analysis:** Review satisfaction scores for each topic to understand user sentiment beyond resolution. Quick resolutions with low scores may indicate a rushed or incomplete experience.
5. **Compare Topic Performance Over Time:** Examine how topic effectiveness evolves over time to prioritize updates or training for underperforming areas.

