



Super Usage Report 2.0 Interpretation Guide

Last updated 2/20/26

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 Super Usage report. To accomplish this, the guide includes two main sections:

Common Findings and Recommendations

This section highlights some of the most common trends and patterns observed across organizations using the Super Usage report. Refer to this section when you want to learn more about a particular data pattern you may be observing in your organization

Detailed Slide Interpretation

This section walks through each slide in the Super Usage report in detail and explains the slide'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 3-4) and Glossary (slide 5) for additional helpful resources

FAQ

Q: Who is included in these numbers?

A: Anyone who is considered an active user (see Glossary for definition). Note that this only measures the licensed (paid) population only – it does NOT measure free Copilot chat (web) usage except for licensed users.

Q: How often should we be refreshing the analysis?

A: In general, we recommend about a cadence of once a month. This is a good balance between capturing the latest data while not being too overwhelming for the analyst. It also allows enough time to transpire between analyses to realistically start to see behavioral changes.

Q: What information do you need to run this report?

A: The report will actually run without organization data, so all you absolutely need is Organization attributed from EntraID. But adding organizational data allows you to do more in-depth cuts

Q: What time frame is each report measured on?

A: Unless otherwise indicated, most charts reflect number of actions per week over the entire time period in the data pull. The time window can usually be changed by adjusting the slider at the top of the chart.

Q: Are there any prerequisites for running this analysis?

A: We recommend at least 1-2 quarters of Copilot usage before running this analysis. That's when you see clear patterns in adoption.

Q: What is the difference between an action and a prompt?

A: Within most surfaces, there are specific actions that are measured – for example, in Word, you can ask Copilot to summarize a document, draft a document, etc. Anything that is NOT mapped to a feature-level category (Create, Analyze, Format) is captured in prompts submitted. In Chat, all actions are categorized as prompts submitted because distinct user action types are not currently being tracked.

FAQ

Q: What is the difference between Intelligent recap and summarize meeting?

A: They are definitely related and both aim to improve meeting productivity. Intelligent recap requires Teams Premium (so if there is a big spike in usage, you might double-check to see whether there was a change in licensing) and requires recording. It also provides a personalized overview and follow-ups, whereas Summarize meeting focuses more on summarizing the key information and highlights from the meeting.

Q: If a meeting is recorded and a summary automatically generated, does that count as an action?

A: You would need to expand the document summary or review the meeting recap. Just generating the summary automatically does not count as an action.

Q: Where do chat actions from Teams get counted?

A: If you initiate a *web-grounded* Copilot Chat request from within Microsoft Teams, it is counted as a Copilot Chat action, not a Teams action. If you initiate a *work-grounded* Copilot Chat request from within Microsoft Teams, it is counted as a Teams action.

Q: What is the difference between usage threshold and usage rank?

A: The Usage Thresholds (Power User tiers) should be your primary thresholds to focus on when tracking progress toward your adoption goals. The Usage Rank (legacy super user tiers) can be a useful way to determine the distribution of usage across your population, highlight the opportunities, and understand the difference in how high vs low users use Copilot.

Q: Why do we use a 9 out of 12 week rolling window?

A: The 9 out of 12 week rolling window is based out of research on habits that it takes an average of 66 days to form a habit. The 3-week gap allows for breaks in the habit, such as when someone is on leave. The literature supports that short breaks in behavior do not necessarily inhibit the habit-forming process.

Glossary

See here for a full description of all Insights metrics and terms:
<https://learn.microsoft.com/en-us/viva/insights/advanced/reference/glossary>
<https://learn.microsoft.com/en-us/viva/insights/advanced/reference/metrics>

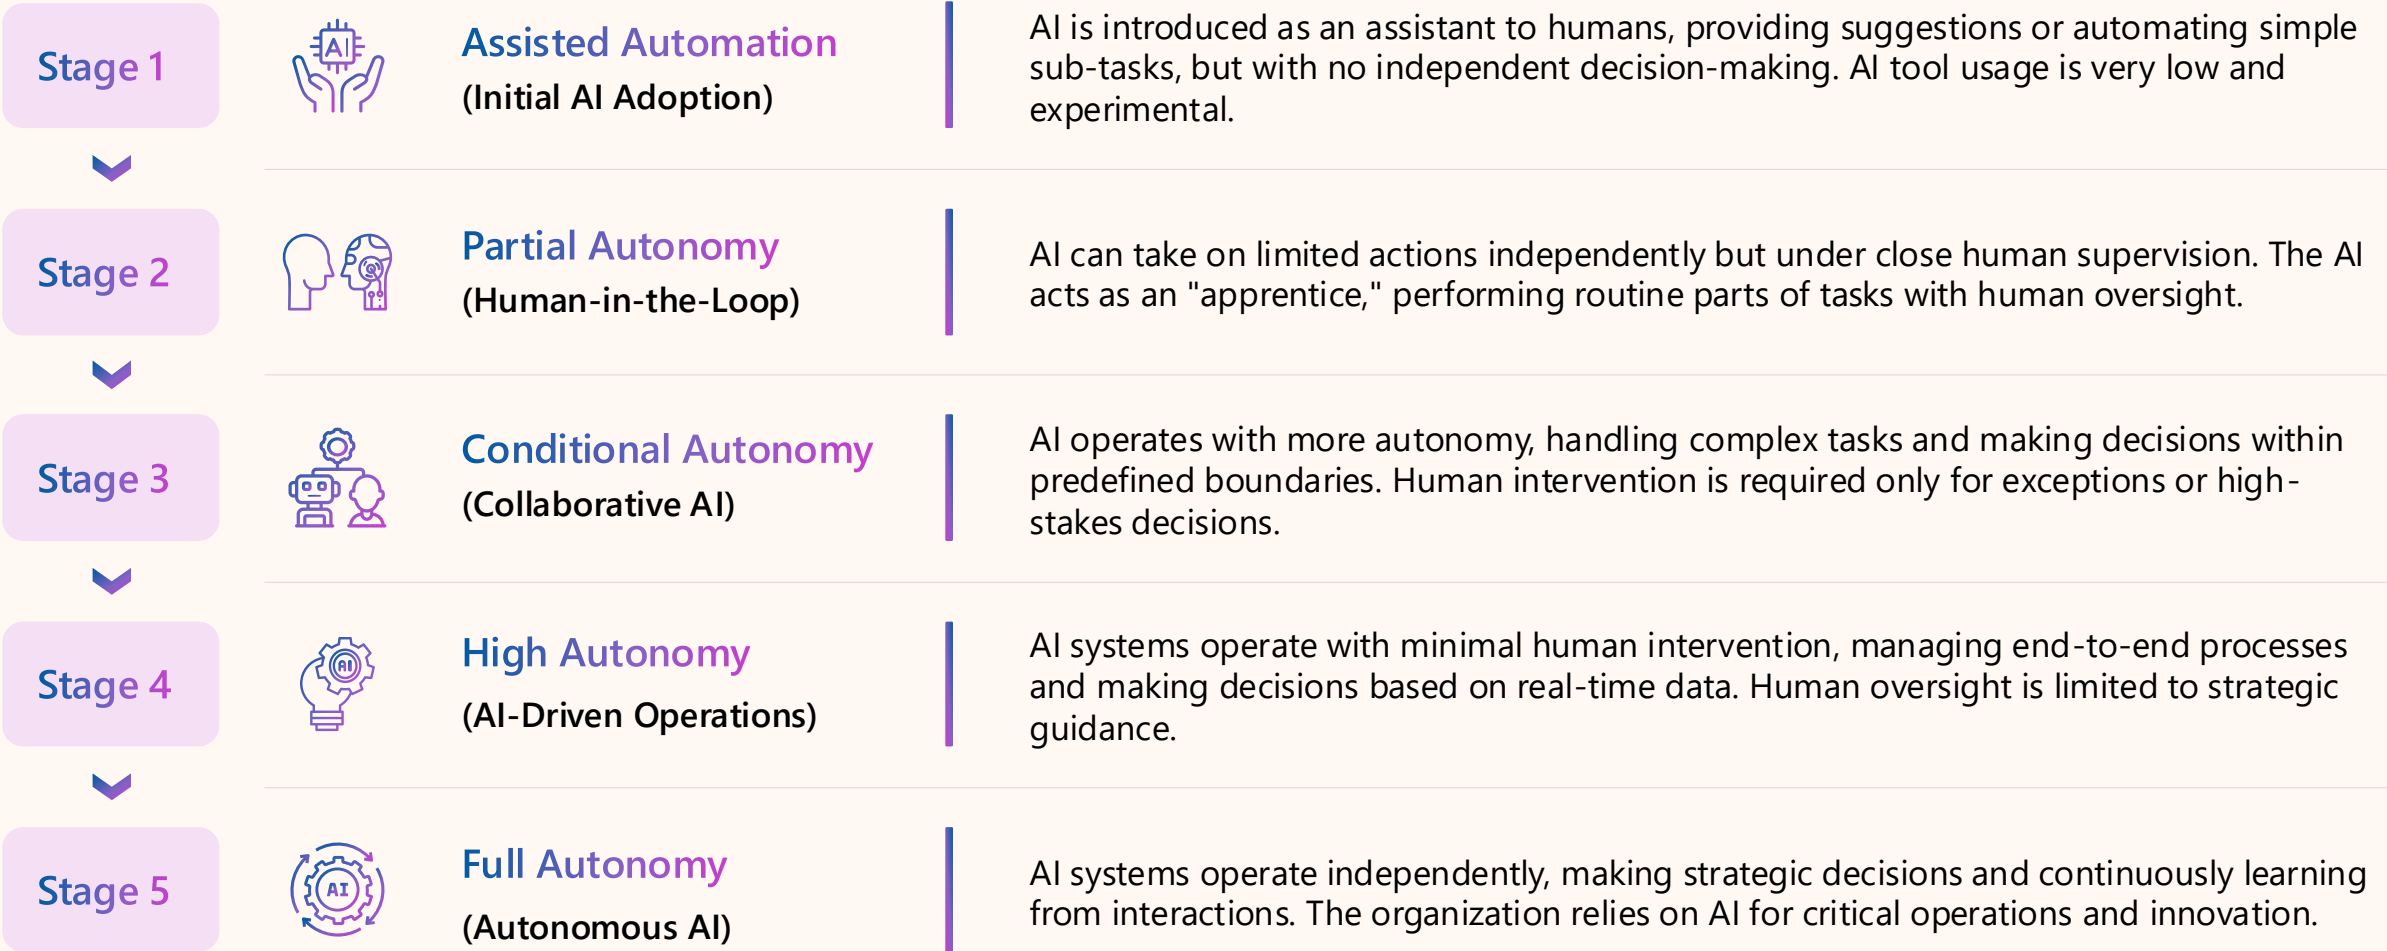
Actions taken	Any Copilot interaction in any surface counts as an action taken
Active connected hours	How much time a person is actively collaborating (at least one meeting or Teams call, email sent, Teams chat sent, or Teams channel message post or reply)
Active days	How many days people are using Copilot in a given time period
Active user	The number of licensed employees who have utilized Copilot during this reporting window.
App breadth	How many surfaces people are using Copilot in a given time period
Assisted hours	<p>Assisted hours are calculated as follows the sum of the following:</p> <ul style="list-style-type: none">- Meetings summarized = length of meeting- Intelligent Recap = 30 mins x # Intelligent Recap actions- All other prompts = 6 mins x # Prompts <p>Note that if the same meeting was summarized and recapped, only the most recent action counts toward assisted hours.</p>
Email hours	Time spent sending and receiving emails. The product does not "read" the contents of the email or log any key strokes.
Enabled days	The number of days when Copilot features are actively enabled and available for use (usually 7 per week). This is different than active days, which reflects how many days Copilot is actually used
Prompts submitted	Within most surfaces, there are specific actions that are measured – for example, in Word, you can ask Copilot to summarize a document, draft a document, etc. Anything that is NOT mapped to a feature-level category (Create, Analyze, Format) is captured in prompts submitted. In Chat, all actions are categorized as prompts submitted because distinct user action types are not currently being tracked.
Surfaces	Used (almost) interchangeably with apps, these are the 9 different places we track Copilot usage for this report: Chat (web), Chat (work), Excel, Loop, OneNote, Outlook, PowerPoint, Teams, Word. Note that we don't have metrics for Loop and OneNote (except days active).
Total Copilot actions	For ease of explainability, we use a custom definition: sum of actions taken in Copilot chat (work), Word, Excel, Outlook, PowerPoint, Teams, and prompts submitted in Copilot chat (web).
Usage tier	<p>Usage tiers are defined based on both volume of Copilot actions as well as the consistency of those actions. See slide 6 for specific definitions of the usage thresholds recommended throughout this report and slides 7-8 for more information about these tiers and how they relate to overall AI maturity.</p> <p>Use usage tiers to benchmark your progress, with "power users" representing those users who use Copilot the most (in terms of actions) as well as most consistently.</p>
Weekly active days	The average number days you took any action with Copilot per week. If a person was active on 2 days during week 1, 3 days during week 2, and 4 days during week 3, the weekly active days during this period would be 3.

The importance of user segmentation

Find ‘**Bright Spots**’ by segmenting users of Copilot based on their usage patterns, specifically focusing on volume (weekly Copilot actions) and consistency (% of weeks with any Copilot usage).

Success measures		Opportunity segment (intervention)		
1	2	3	4	5
Power Users <p>Power Users represent the ideal user who maximizes the potential of Copilot and are both consistent and high-volume users.</p> <p>This group is likely to be a minority of users but can be seen as aspirational group for deploying an AI adoption strategy.</p> <p>Averaging 20+ weekly Copilot actions and any use of Copilot in at least 9 out of past 12 weeks</p>	Habitual Users <p>Habitual Users are consistent users of Copilot, with lower volume than Power Users.</p> <p>They represent ‘everyday’ users who have successfully adopted Copilot into their routine.</p> <p>Averaging 8+ weekly Copilot actions and any use of Copilot in at least 9 out of past 12 weeks</p>	Novice Users <p>Novice Users are users with potential to become Habitual Users, but who may need additional support to avoid lapsing into lower usage.</p> <p>Averaging at least one Copilot action over the last 12 weeks</p>	Low Users <p>Low Users are either early in their adoption journey or require significant assistance with onboarding and utilizing Copilot.</p> <p>Having any Copilot action in the past 12 weeks</p>	Non-users <p>Non-users are individuals who are enabled on Copilot, but do not use it.</p> <p>Zero Copilot actions in the last 12 weeks</p>

Overview of the 5 Stages of AI Maturity



User Segments in Each Stage

		Power User	Habitual Users	Novice Users	Low Users	Non-users
Stage 1	Assisted Automation	Practically none. Maybe a couple of pioneers. A true Power User consistently uses AI 20+ times/week over months, but at Stage 1, almost no one does. Even using AI ~5 times/week would make someone stand out.	Very few. A handful of early adopters may use AI weekly, meeting the 9-of-12-weeks consistency threshold. Volume is low, but use is regular.	Some. Novices at Stage 1 use AI sporadically – averaging ~1 action/week over a quarter, but inconsistently. Many tried it during onboarding; ~20–30% may fall in this group.	Many. Most Stage 1 users try AI once or twice, then stop – often due to unclear value. Converting them is a key opportunity.	Majority. Most haven't tried AI yet. Awareness, relevance, or time are key barriers. Reducing this is critical.
Stage 2	Partial Autonomy	Emerging. At Stage 2, a small but growing group (2–5%) become Power Users – those using AI heavily (tens of actions/week), often due to automation-friendly roles.	Expanding cohort. At Stage 2, 10–20% of employees use AI weekly (9 of 12 weeks), showing a shift from novice to habitual use.	Still significant. At Stage 2, ~30–50% are novices using AI irregularly – down from Stage 1 as more become habitual users, though some remain low or inactive.	Shrinking. At Stage 2, low users drop to ~10–15% as some move up to novices, though others remain on the fence.	Declining. As adoption grows, non-users drop to ~20–30%. By Stage 2, over half of enabled users have tried AI at least once.
Stage 3	Conditional Autonomy	Growing presence. The number of Power Users continues to increase as more users become proficient with AI tools and integrate them into their daily workflows.	Increasingly common. A significant portion of users now regularly use AI tools, forming a stable base of habitual users.	Decreasing. As users gain experience, fewer remain novices, shifting into habitual or power user roles.	Minimal. Few low users, as most users have either increased their usage or stopped using the tools altogether.	Rare. Very few users remain who have not engaged with AI tools at all.
Stage 4	High Autonomy	Significant proportion. A large number of users are now power users, heavily relying on AI tools for their daily tasks.	Majority of users. Most users are now either habitual or power users, regularly using AI tools in their workflows.	Few. The number of novice users is now quite small, as most users have become proficient with AI tools.	Very few. Most have either scaled up or stopped, leaving minimal low users.	Almost none. Very few users remain who have not engaged with AI tools at all.
Stage 5	Full Autonomy	Predominant. The majority of users are now power users, heavily relying on AI tools for their daily tasks.	Nearly all users. Most users are now either habitual or power users, regularly using AI tools in their workflows.	Rare. The number of novice users is now very small, as most users have become proficient with AI tools.	Almost none. Most users have either scaled up or dropped off, leaving very few low users.	Virtually none. Very few users remain who have not engaged with AI tools at all.



Common Findings and Recommendations



Summary

This section highlights some of the most common findings in the Super Usage report and explains what is typical, how you might interpret such findings at your organization, and suggests next steps for action. A quick summary of commonly-found trends are listed below, with additional details on the following slides.

	Finding	Recommendation
1	Usage will be uneven – high intensity usage and pockets of low usage.	<p>The high intensity usage is proof that Copilot works</p> <ul style="list-style-type: none">• Learn as much as you can about your super users. What was their adoption journey? What are their work patterns? How is Copilot impacting their work?• Understanding their specific journey and milestones enables us to replicate and scale adoption across the company
2	11–12 weekly actions leads to habit formation; 20 weekly actions begins to differentiate power users	<p>Help users reach these thresholds quickly</p> <ul style="list-style-type: none">• Prompt-a-thons for hands-on experience• Use Copilot Chat (web) for daily tasks- this habit extends into other surfaces
3	Teams with super user managers have more super users	<p>Role modeling matters</p> <ul style="list-style-type: none">• Consider targeted enablement for managers• Leaders should use Copilot visibly and share real use cases with teams
4	High usage is sticky, low usage is also sticky	<p>Adoption takes time; delaying license allocations further pushes out AI readiness</p> <ul style="list-style-type: none">• Identify groups that may need additional focus• Remove barriers (e.g., transcription for Teams, web grounding for Chat)
5	Most users start their journey in Copilot Chat and Teams	<p>Monitor app usage over time</p> <ul style="list-style-type: none">• Usage typically spreads to additional surfaces as users gain confidence• Review super users’ adoption journey to see what apps resonate with them
6	Changes in work patterns based on Copilot usage will vary by role	<p>Certain teams may naturally gravitate toward certain use cases for Copilot</p> <ul style="list-style-type: none">• Identify teams who have heavier usage of certain apps, and learn more about what they are doing• Use this information to provide specific, tangible ideas on how others can tackle similar challenges

General usage trends

Sudden spike in usage



What is typical

- 1 It's not uncommon to see a plateau for a few months and then a sudden spike. Usually spikes happen with new company announcements, leadership messaging, or removal of blockers (e.g., transcription enabled).

Recommendation for action

Consider what factors could be driving the increase. Was there a big leadership push to use Copilot? Prompt-a-thon? Were a significant number of people were given access to Teams chat or another feature or app?

Look at specific app and feature usage trends to see what might be driving the increase.

General usage trends

Sudden drop in usage



What is typical

- 1 If you add a large number of licenses in one month, you might see a dip in usage as the new licensees ramp up their usage.


We frequently see a dip in December (and into January) related to end of year time off; in Europe, we often see the same thing for summer holiday in August. Consider if there are other holiday or business-related slowdowns or shutdowns that might be impacting usage.

Recommendation for action

Look at specific app and feature usage trends to see if these drops are across the board, or if they might be related to a specific change in access (for example, were certain licenses now restricted from high-use features like transcription in Teams meetings?)

General usage trends

Low usage tiers are stagnant



What is typical

1

Ideally we would see all usage tiers increasing their average actions over time, but it is not unusual to see those in very low usage tiers remain below the habitual usage threshold. Have these patterns remained steady over several months? Are you observing any changes in usage of specific apps or features?

Recommendation for action

- Remember that adoption is a curve and not a switch. It takes time to go through the adoption journey. In your data, look at the percentage that are above the habit threshold and the rate it is growing every month. These super users are proof that Copilot works. Now we can learn from them to accelerate adoption for others.
- Use the adoption habits and patterns of Super Users as a path you can share for others to emulate. Look at their weeks since activation: How long did it take them to progress toward being a Super User? What apps and features did they adopt more quickly?
- In general, we recommend continuing to expand and enable others in your organization, even if a significant percentage has yet to become super users. If you delay enabling others with Copilot and wait till you hit 100% habit formation in the existing population, then you are pushing back AI readiness for your company even further.

App or feature usage trends

Usage is concentrated to one or two apps



What is typical

- 1 It is not unusual to see Copilot Chat and Teams features most frequently used, as Copilot Chat and Meeting Summarization are common entry points for users.

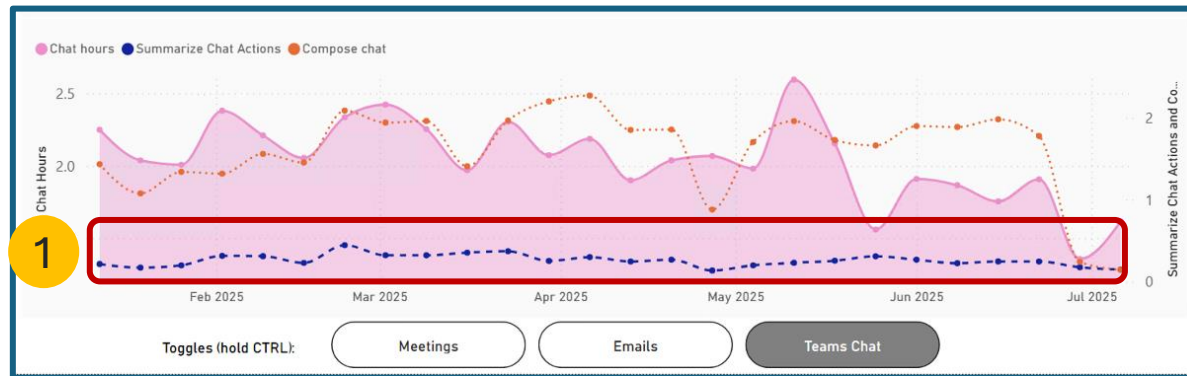
Ideally over time we would see increases across all apps as users gain confidence in using Copilot.

Recommendation for action

- If we see increases in some apps while others remain flat, there may be an opportunity to upskill or communicate the benefits of Copilot in these lesser-used surfaces.
- Consider how specific teams with similar work profiles are utilizing various apps; if a team does a significant amount of documentation work, you should expect to see a rise in Copilot in Word features. This can guide enablement efforts toward more targeted examples.
- Observe if this pattern differs for Super Users by applying a filter on Usage Tier. If Super Users are also limiting their Copilot usage to a relatively small number of apps, consider whether a broader enablement effort is required to better inform employees about the benefits of specific apps and features.

App or feature usage trends

Some features are not being utilized



What is typical

- 1 We do see some features typically get adopted earlier than others, particularly chat. We also see Intelligent Recap usage increase quickly once enabled. Other features, however, will often vary by role; customer service roles may see higher usage of Outlook features, whereas data analyst roles may see higher usage of Excel features. Look at results for different organizational groups to see how their usage may differ.

Recommendation for action

- Comparing super users vs. everyone else may surface some structural barriers. Are super users using this feature while the rest of the company doesn't? This might indicate that there may be policies restricting usage of certain features.
- If you see some features with little to no use, consider whether a targeted effort to publicize the value of those features would be beneficial. Create communication and enablement plans that highlight specific use cases for underutilized features to illustrate real-life examples of how these features can address actual needs at your organization.

Trends across groups

One team's usage is significantly higher/lower than others

1

Region	Users	Weekly Copilot Actions
Japan	19	62.1
ASIA	60	40.9
Singapore	13	37.6
Australia	14	37.1
India	27	36.9
Canada	10	36.4
United States	438	30.9
Ireland	76	28.8
United Kingdom	48	28.7
EMEA	72	28.3

What is typical

1

It is not unusual to see usage differences across teams. Look at the specific populations to determine if there are specific licensing differences that might be driving the discrepancy. Which teams are fully licensed versus partially licensed? How long have team members had licenses?

In addition, remember to look at group N size, as smaller groups often show greater volatility (reporting either significantly higher or lower usage).


Recommendation for action

- Identify groups where there is high density of Super Users. These groups have clearly embedded Copilot into their work flows, and it is likely that you will find sources of best practices from these teams.
- Identify teams with an unusually high or low usage of a particular app. The nature of each team's work may impact Copilot usage patterns. Teams with higher than average use of certain apps or features may be able to provide recommendations and serve as inspiration for other teams.

Work patterns

Changes in email usage

1



What is typical

1

We often see a decrease in email time as use of Copilot increases and employees are able to be more efficient with their email OR are shifting more to Teams Chat.

However, in other organizations and/or roles, we may also actually see an increase in email, particularly in email-heavy roles such as customer service.

Recommendation for action

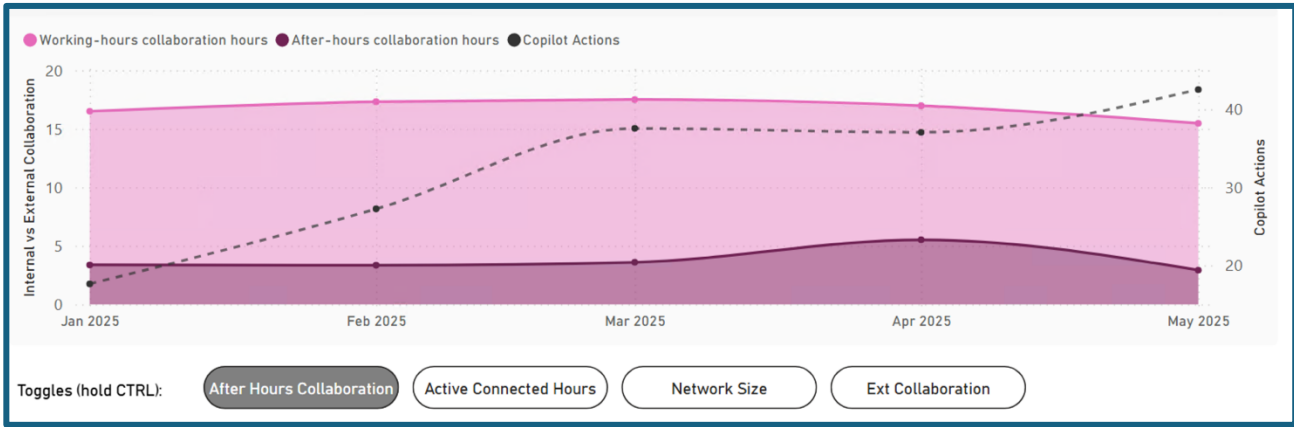
- In both cases, look at the populations with the greatest shifts in email habits. What types of roles are impacted most? Do we see corresponding changes in other app usage, such as Chat, that might be related to the change in email use?

Consider holding focus groups with these teams to learn more about specifically how they are using Copilot to improve their email productivity – and share these findings with other teams in similar roles.

Work patterns

Not seeing an impact on work patterns

1



What is typical

- 1 Note that impact on work patterns will vary by customer and roles. We frequently see patterns including after-hours work declining with increased Copilot usage and increased Copilot usage for those with higher collaboration hours, but your results may differ based on who is using Copilot and how they are using it.
- It is important to remember on these types of charts that we are showing correlation, not causation.


Recommendation for action

- Try analyzing collaboration patterns for different usage cohorts over time - are super users seeing a greater impact on their work patterns? You may see that the employees who have the greatest potential for Copilot to make a meaningful impact on their work (i.e., those with highest workload, most meetings, etc.) are indeed the ones using it most.
- Is this seasonal? Or are there shifts in how people collaborate? Are they attending fewer meetings but accepting more to follow along with Copilot notes? Look at this data in conjunction with other work patterns to get a fuller picture of how they may be changing.

Work patterns

Changes in collaboration hours

1



What is typical

1

Collaboration hours might rise with increasing Copilot usage. This may be due to employees accepting a greater number of meetings as their ability to catch-up via recap increases. Especially in customer-facing roles, it is not unusual for increased Copilot actions to coincide with an increase in external collaboration, as Copilot is likely enabling them to work more efficiently and serve more customers.


However, other customers may see decreases in after-hours collaboration hours as users are able to collaborate more efficiently through Copilot's facilitation of asynchronous work, reducing the need for after-hours meetings.

Recommendation for action

- Analyze the pattern across different types of roles to see how this impact may differ; look at super users versus all others as well.
- Look at associated work patterns as well: are you also seeing changes in after-hours work? Meetings and email hours? Try to evaluate the work patterns holistically to determine what the true impact on work might be.
- If available, review any Copilot sentiment data for these groups as well. Are they reporting any different patterns in their own estimations of Copilot's impact?

Work patterns

Shrinking internal network size



What is typical

1

This decrease in network size is common and could be due to a greater number of people using Copilot chat for search, rather than reaching out to internal colleagues – thus shrinking the internal network size.

Recommendation for action

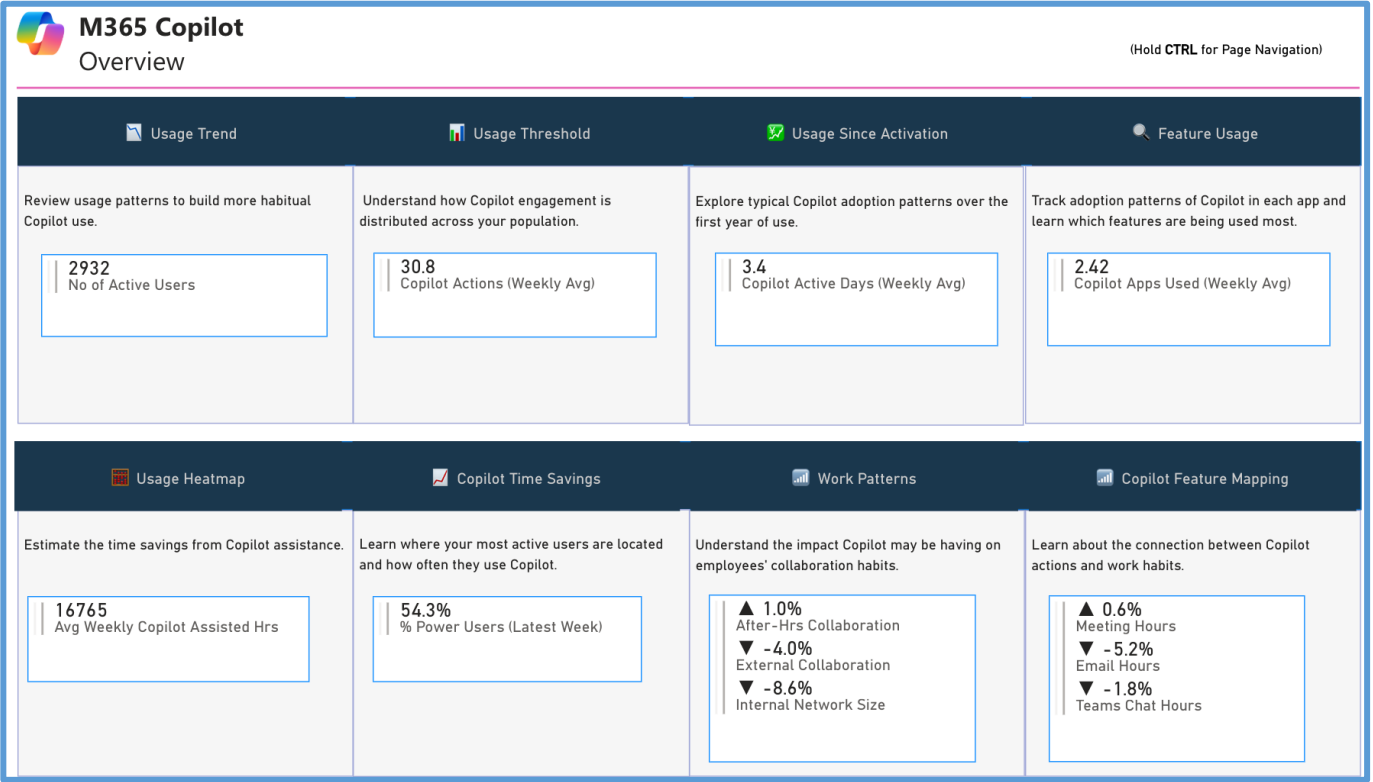
- Analyze the pattern across different types of roles to see how this impact may differ; look at super users versus all others as well.
- Look at associated work patterns as well: are you also seeing changes in after-hours work? Meetings and email hours? Try to evaluate the work patterns holistically to determine what the true impact on work might be.
- If available, review any Copilot sentiment data for these groups as well. Are they reporting any different patterns in their own estimations of Copilot's impact?

Detailed Slide Interpretation



This page is designed to provide a high-level look at the subsequent report sections.

Each section also includes one headline statistic; click on the section header to navigate directly to that page to learn more.



Usage Summary

1 What are the overall usage metrics for this group?

INTERPRETATION:

- **Enabled users:** The number of licensed employees in this group. Use this information along with headcount to identify where licenses are currently allocated and where you may consider additional licenses for groups that are low relative to other teams.
- **Active users:** The number of licensed employees who have utilized Copilot during this reporting window. We ideally want this number to be as close to 100% as possible, which would indicate full utilization of the allocated licenses. Where it is lower than 100%, increasing numbers week-over-week is a positive indicator of progress.
- **Weekly Copilot actions:** The average number of Copilot actions each user is completing on a weekly basis
- **Weekly change rate:** The average rate of change. Use this to assess how the rate of usage intensity has changed over this time period. Ideally this number would be positive, indicating positive growth week-over-week.
- **Power users:** This number provides an indication of usage intensity within a team. See the definitions in the glossary section for further information on how this threshold is defined. Use this to gauge users' adoption and habit formation over time.

2 How many Copilot actions are users averaging each week?

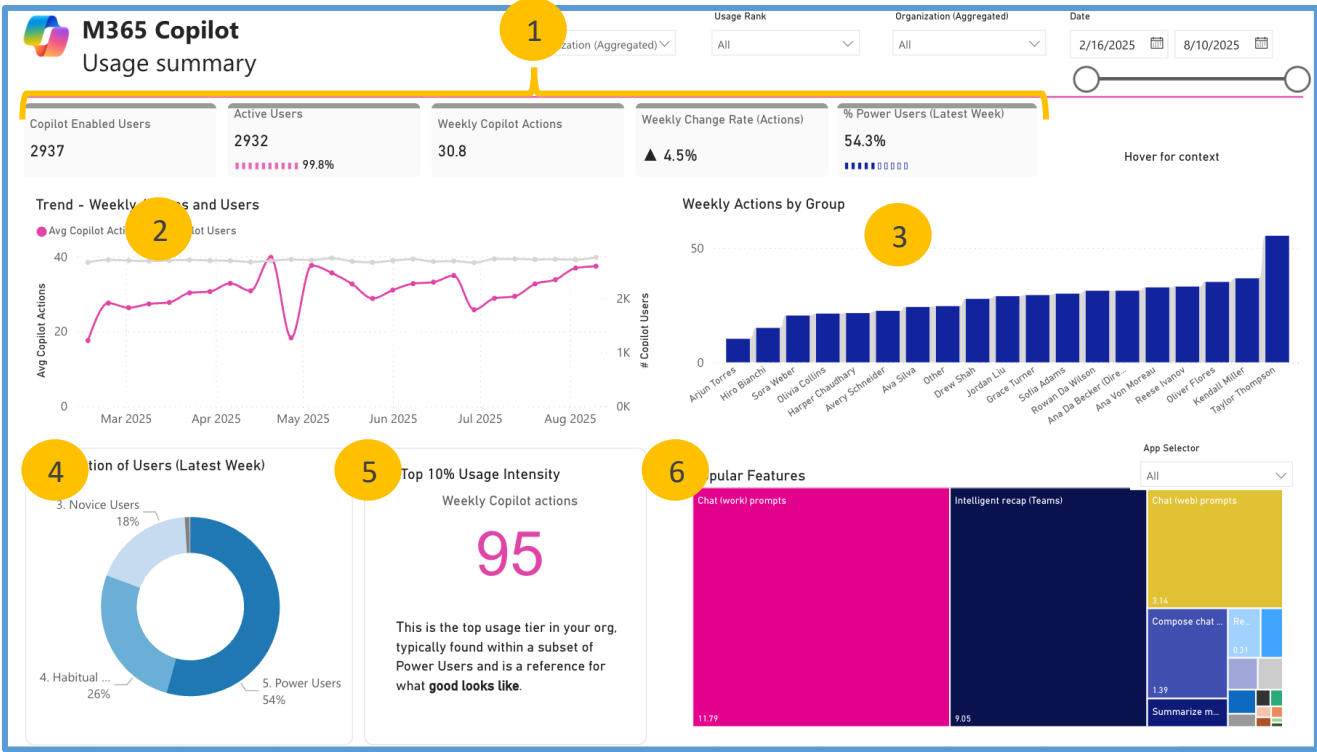
GUIDANCE:

- Upward trend shows that people are completing more Copilot actions each month. We want to see this “up and to the right” trend continuing
- Note that these are AVERAGE numbers, not total actions. If you add a large number of licenses in one month, you might see a dip in usage as new licensees ramp up usage
- Similarly, if you see a big spike, you might consider what factors could be driving the increase. Was there a big leadership push to use Copilot? Prompt-a-thon? Etc.

3 How does Copilot usage vary across teams?

From this data, we see that Taylor Thompson’s group has the most Copilot usage per user, while Arjun Torres’s group has the lowest.

INTERPRETATION: The average weekly Copilot actions for each team. Consider how long certain teams have had their Copilot licenses and how that might affect their adoption rate; teams with longer Copilot usage are likely to see higher usage, whereas teams new to Copilot may see lower usage as they build their Copilot capability.



4 What is the usage tier distribution?

From this data, we see that 54% of this group were considered power users last week.

INTERPRETATION: See the definitions in the glossary for information on how these thresholds are defined.

5 How much are our top users using Copilot?

From this data, we see that our top 10% are completing 95 actions per week on average.

INTERPRETATION: Use this number to define what usage looks like for top users; this can be a good reference point when setting Copilot goals for the organization.

6 Which actions within each surface are people completing most often?

Copilot chat (work) prompts submitted is the #1 action taken (11.79 actions per week on average), followed by Intelligent recap in Teams (9.05). This indicates that Copilot Chat is by far the surface with most actions.

INTERPRETATION: Note that the colors on this chart correspond with each app– dark blue represents all Teams actions, light blue represents all Outlook actions, etc. Note that for Copilot chat (both web and work), user level actions are not categorized, so all actions are shown as “prompts submitted.”

Usage Trend

1 What are the overall usage metrics for this group?

INTERPRETATION: Use these metrics to monitor the overall Copilot usage for your group. Click on the buttons at the bottom of the chart in more detail.

2 What percent of licensed users are active?

GUIDANCE: Use this to track the percentage of licensed users who are actively using Copilot.

3 How many Copilot actions are users averaging each week?

- GUIDANCE:
- Upward trend shows that people are completing more Copilot actions each month. We want to see an “up and to the right” trend, but it’s not uncommon to see a plateau for a few months and then a sudden spike.
 - Note that these are AVERAGE numbers, not total actions. If you add a large number of licenses in one month, you might see a dip in usage as new licensees ramp up usage
 - Similarly, if you see a big spike, you might consider what factors could be driving the increase. Usually spikes happen with new company announcements, leadership messaging, or removal of blockers (e.g., transcription enabled).

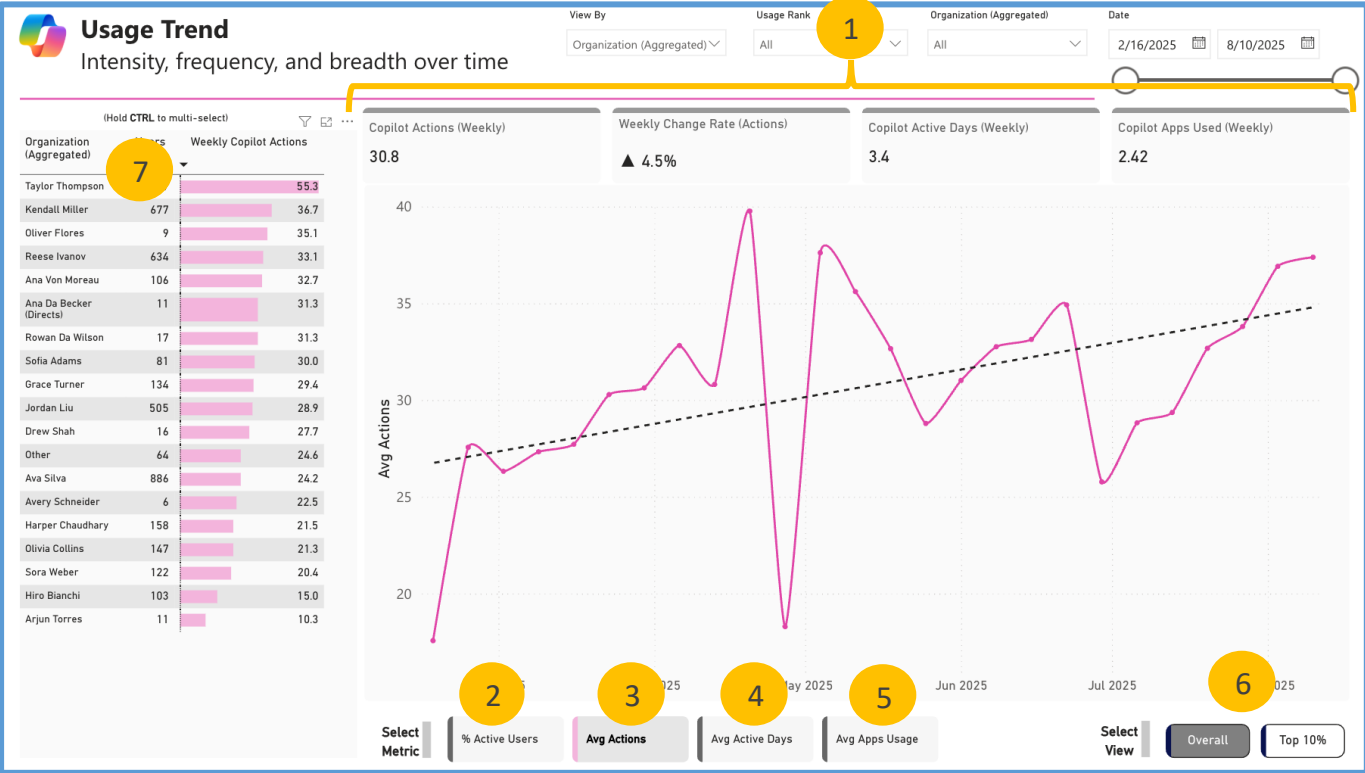
4 How many days are users active on Copilot each week?

GUIDANCE: We want this number to as high as possible, with an eventual target of 4+ days per week. Higher numbers here indicate more habitual, daily usage.

5 How many apps are users using each week?

INTERPRETATION: We want to see app breadth growing. When this is happening, it tells us that Copilot usage is increasing across apps – so we are building stronger Copilot habits in multiple surfaces, not just in 1. It also signals that Copilot is getting embedded into workflows and isn’t just an experimental tool.

GUIDANCE: Don’t worry too much about attaining very high app breadth (5+ apps), as this will generally be rare. Focus more on the mid-range (3-4), as this is likely the better indicator of true growth in app breadth.



6 View overall results or top 10% view?

RECOMMENDATION: Toggle from Overall to Top 10% to switch your view between all results and a view that shows the Top 10% trend compared to all other users. This view is helpful to understand the difference in usage trend for top users – and what patterns you might aspire to for other users.

7 What is the average Copilot usage for different teams?

RECOMMENDATION: Click on individual teams to view their Copilot usage trends over time. If these numbers are significantly different, you might look more closely into the specific populations to determine what might be underlying the discrepancy (e.g., time since licensing, which groups have licenses, etc.)

Usage Thresholds

1

What percentage of users belong to different usage tiers, and how has this changed over time?

INTERPRETATION: As we see Copilot gain traction in an organization, we expect to see an increase in power and habitual users, and a decrease in low and non-users. Such a shift in usage categories tells us that Copilot usage is becoming more integrated into people's daily work, establishing Copilot as a routine part of their ways of working.

GUIDANCE: Common trends may include spikes (for example, when there is a large-scale enablement effort) or dips (for example, when a large group of new licenses are issued and those users are still ramping up their usage).

2

How do these percentages vary across organizational groups?


INTERPRETATION: As with the view over time, we want to see larger percentages of users within each team in the habitual power user categories, and lower percentages in low or non-user categories.

RECOMMENDATION: When looking across teams, look at groups with a high percentage of power users – what can we learn from them? Might they be a good source of best practices for other teams? Focus also on teams with few power users, as they might be a good target for additional enablement.

3

When should I use usage threshold versus usage rank?

RECOMMENDATION: The Usage Thresholds (Power User tiers) should be your primary thresholds to focus on when tracking progress toward your adoption goals. The Usage Rank (legacy super user tiers) can be a useful way to determine the distribution of usage across your population, highlight the opportunities, and understand the difference in how high vs low users use Copilot.



Usage Heatmap

Important notes on reading this slide:

- 1) Teams with less than 5 employees will not have results shown.
- 2) Where shading is applied, shades of green indicate higher numbers and shades of red indicate lower numbers, with yellow and orange shades in the middle. Use these colors to quickly assess comparisons across groups and across time periods.

1 What are the overall usage metrics for this group?

INTERPRETATION: Use these metrics to monitor the overall Copilot usage for your group. Click on the buttons to the left of the heat map to view each of these metrics over time with more detail.

2 How has the number of enabled users per group changed over time?

INTERPRETATION: This number represents the number of licensed employees per group by week.

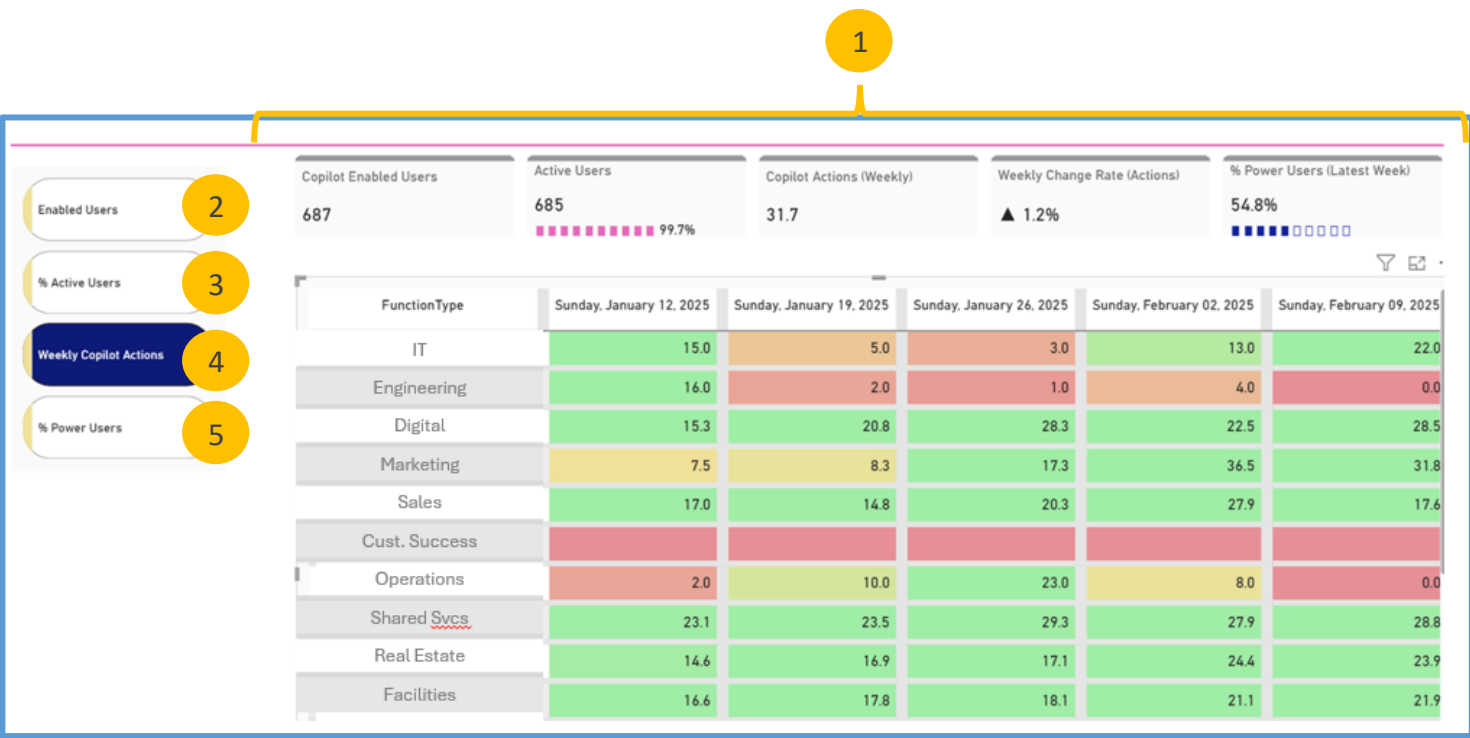
RECOMMENDATION: Use this information in comparison with group headcount to identify where Copilot licenses are currently allocated, and where you may consider additional licenses for groups that are low relative to other teams.

3 Of licensed users, how many are actually using Copilot?

INTERPRETATION: This number represents the number of licensed employees per group by week who have utilized Copilot during this reporting window.

GUIDANCE: We ideally want this number to be as close to 100% as possible, which would indicate full utilization of the allocated licenses. Where it is lower than 100%, increasing numbers week-over-week is a positive indicator of progress.

RECOMMENDATION: Focus on groups with lower percentages (shaded in red) to identify teams where additional enablement – or potential license reallocation – may be warranted.



4 How many Copilot actions are teams averaging per week?

INTERPRETATION: This number provides a closer look at where team members are utilizing Copilot the most.

GUIDANCE: As with active users, increasing numbers week-over-week is a positive indicator of growing interest and utilization of Copilot.

RECOMMENDATION: Focus on groups with lower averages (shaded in red) to identify teams where additional enablement – or potential license reallocation – may be warranted.

5 What percentage of each team are considered power users?

INTERPRETATION: This number provides an indication of usage intensity within a team. See the definitions in the glossary section for further information on how this threshold is defined.

RECOMMENDATION: Look at groups with a high percentage of power users – what can we learn from them? Might they be a good source of best practices for other teams? Focus also on teams with few power users, as they might be a good target for additional enablement.

Usage Since Activation

1

What does a typical Copilot adoption curve look like?

Unlike the previous pages, this page looks at time series based on activation date. Notice that x-axis here is week since activation; this means we are observing the trend for each user's week over week adoption trend, regardless of when their license was enabled. From this data, we see a steady increase over the first 10-15 or so weeks of licensure, after which usage stabilizes somewhat.

GUIDANCE: Typically, we will see it take around a minimum of one quarter of usage (12 weeks) for users to start becoming habitual Copilot users. Remember that high usage becomes sticky quite quickly – but low usage is similarly sticky. If your organization is showing a longer journey to habitual use, you might consider targeted enablement to help users realize the value of Copilot more quickly.


RECOMMENDATION: Toggle to “By App” **2** to view the adoption pattern by individual app or “By Usage Rank” to view the adoption pattern by each usage tier.

3

How quickly are different teams adopting Copilot?

From this data, we see that Japan has the most rapid increase in their Copilot use over this time period (average +3% per week).

INTERPRETATION: These numbers represent the average week-over-week increase in Copilot actions for each team listed here. As you compare across teams, consider how long certain teams have had their Copilot licenses and how that might affect their adoption rate; teams with longer Copilot usage are likely to see more of a plateau in their actions, whereas teams new to Copilot are likely seeing a more rapid increase.



Feature Usage

1

Which actions are people completing most often?

Copilot chat (work) prompts submitted is the #1 action taken, followed by Intelligent recap in Teams. Copilot Chat (Work) invoked from Teams is the third most action. This indicates that Copilot Chat is by far the surface with most actions.

INTERPRETATION: Note that the colors on this chart correspond with the colors on the above chart – dark blue represents all Teams actions, light blue represents all Outlook actions, etc. Note that for Copilot chat (both web and work), user level actions are not categorized, so all actions are shown as “prompts submitted.”

RECOMMENDATION: Toggle by each app

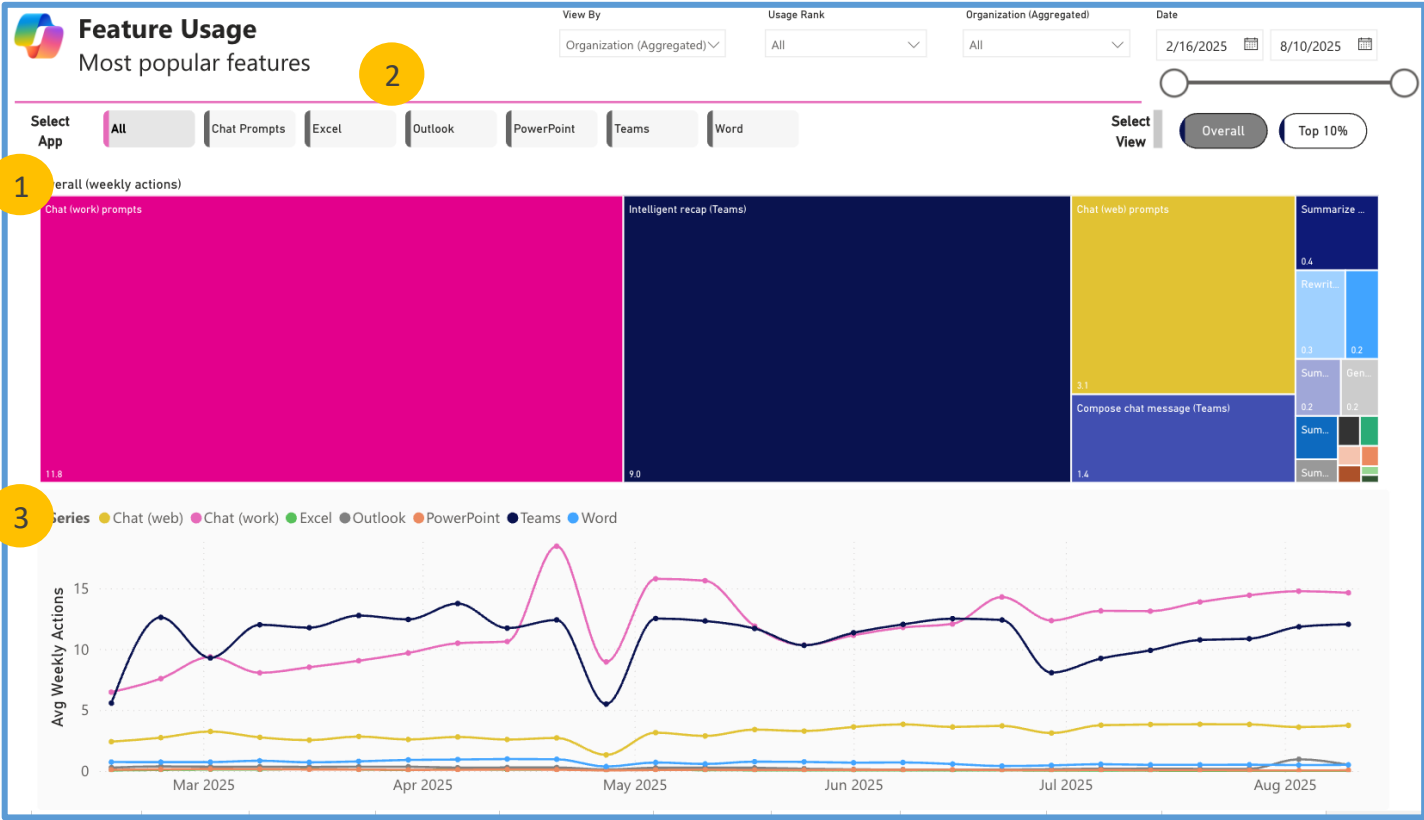
2

to view the feature usage pattern by individual app.

3

How has weekly app usage changed over time?

For this data, we see Teams and Chat (work) have consistently been the top apps.



RECOMMENDATION: Ideally, over time we see increases across apps and features as users gain confidence in using Copilot. Adjust the time horizon to see how the numbers change over time; since they reflect just the highlighted time window, if you adjust to include more recent usage, you should see these numbers increase. If we see increases in some apps and features while others remain flat, there may be an opportunity to upskill or communicate the benefits of Copilot in these lesser-used surfaces.

These charts can also be a useful way to compare and contrast how teams with similar work profiles are utilizing various apps. If a team does a significant amount of documentation work, you should expect to see a rise in Copilot in Word features. This can guide enablement efforts.

BONUS TIP: Comparing power users vs. everyone else can indicate some structural barriers. Are power users using Intelligent Recap while the rest of the company doesn't? This might indicate that there may be policies restricting usage of certain features.

Group Scatterplot

1

How do different teams compare in terms of overall frequency and intensity of Copilot usage?

From this data, we see Taylor Thompson’s team used Copilot both the most days as well as most frequently.

INTERPRETATION: The size of each bubble reflects the relative size of the team – so larger bubbles represent larger teams. Look for teams that are outliers, either falling below most other teams or outpacing them. Pay particular attention to large teams that fall into either of these categories, as they represent a particularly large segment of the population that could be a good target for either best practices or further enablement.

2

How do different teams compare in terms of frequency and intensity of usage of specific apps?


RECOMMENDATION: Toggle between the different apps to see how relative placement on the chart changes. Do certain teams move significantly relative to their overall position when you toggle to a particular app? This could be an opportunity to learn more about what that team is doing and see how it could be applied across other teams.

3

How much are well-connected teams using Copilot?

INTERPRETATION: This chart plots teams based on their network size and percentage of power users.

RECOMMENDATION: Use this chart to identify teams that are particularly influential (i.e., those with large networks) who also have a high percentage of power users. These teams might be good targets for sources of best practices and role modeling for others. Similarly, if there is a well-connected team with an unusually low percentage of power users, they might represent a good opportunity for targeted enablement.



Copilot Assisted Value

1 How much time does Copilot assist with – overall, and by user? And what does that translate to in terms of actual value?

INTERPRETATION: The first two numbers represent the total value of the Copilot actions taken, first for the entire population and then on average per user. These numbers are then translated to a monetary value using the multipliers in section 2.

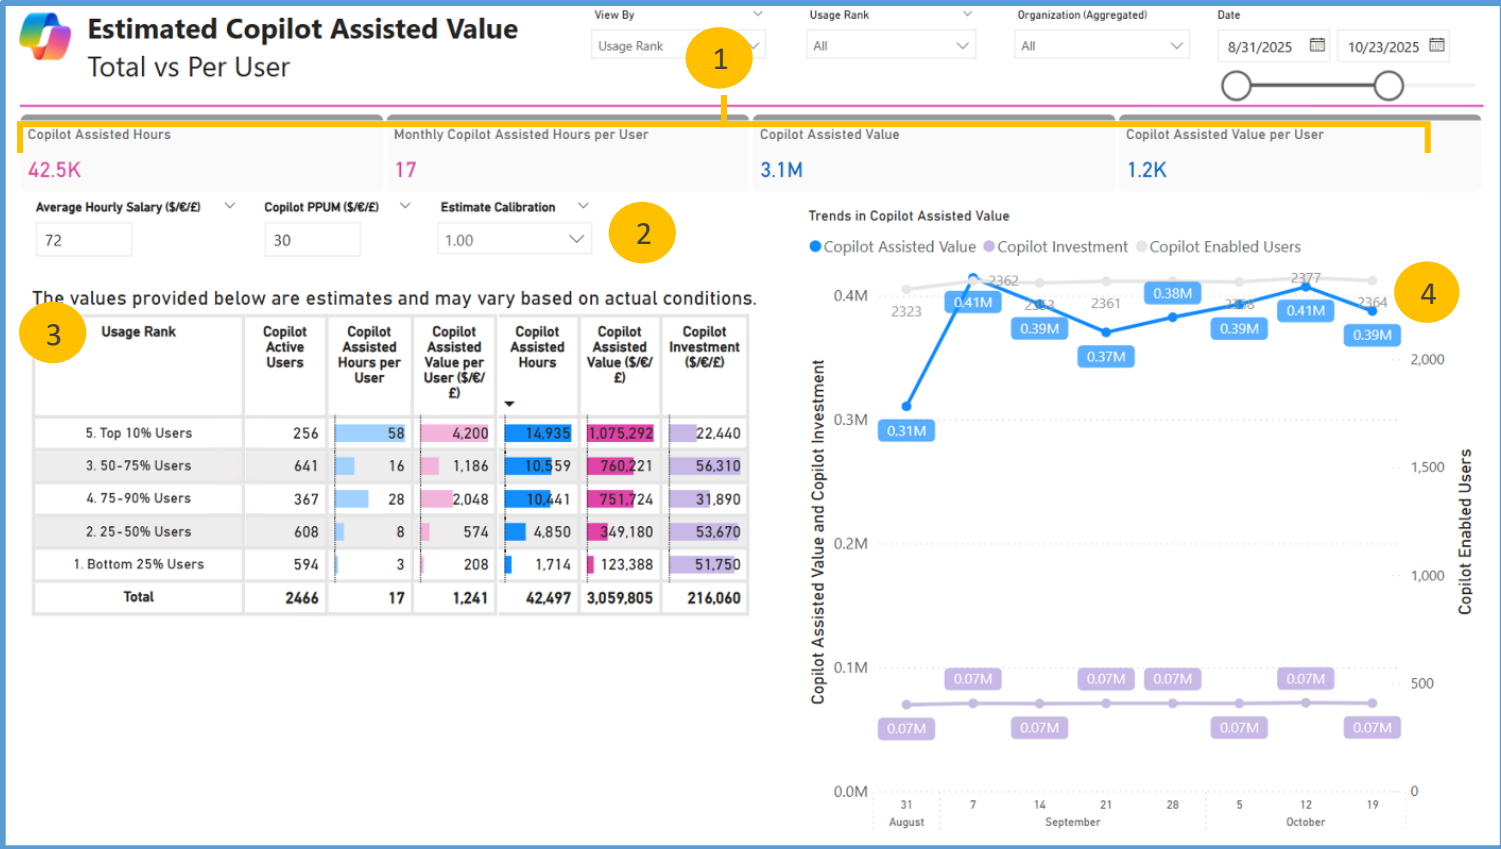
2 How can I customize the estimates/assumptions?

GUIDANCE: By default, the template will populate with certain values for average hourly salary and Copilot PPUM. If you believe these are too high or low for your situation, they can be adjusted in this section. Similarly, if you want to adjust how conservative the estimate is for assisted hours, you can adjust the estimate calibration.

3 How do usage and value differ across usage tiers?

INTERPRETATION: Compare differences in usage and value first at a per user level; in this data, the top 10% of users average 58 hours of Copilot assisted time, compared to 28 hours per user for those in the top 75-90%. This tells us that the top 10% are getting more than twice as much value from their Copilot usage as the next 15%. However, the value both are deriving far exceeds the \$30 license cost, indicating very strong ROI for these groups.

We can do the same for the total populations by usage tier; rather than averages per user, these figures represent the entire aggregate hours and value for the group, compared to the total Copilot investment for that group. In this data, we see that our bottom 25% of users have a total of 1,714 assisted hours, valued at \$123,388 of assisted value; still, when compared to the total cost of \$51,750 for their licenses, even our bottom 25% of users are still seeing a positive ROI.



4 How have assisted value, Copilot investment, and enabled users trended over time?

INTERPRETATION: Copilot investment (purple line) reflects the cost of all licenses, and enabled users (gray line) reflects the number of allocated licenses; assuming all (or almost all) licenses have been enabled, these lines will generally run parallel to one another. The assisted value line in blue charts the monetary value of Copilot usage for each week; so, as Copilot usage increases, the assisted value line will also increase.

GUIDANCE: Early after adopting Copilot, a primary focus will often be on ensuring the assisted value exceeds the investment cost, indicating positive ROI. Once adoption is strong, the goal will be to see a steady increase in assisted value, which indicates that ROI is also increasing over time, since the investment itself holds steady (unless there is a change in the number of licenses purchased).

Work Patterns

1 How has Copilot impacted work patterns?

For this data, after-hours collaboration and active connected hours have both decreased over this time period, whereas internal network size and external collaboration hours have increased somewhat.

2 What is the average Copilot usage for different teams?

For this data, we see IT has the highest usage, versus Legal which has the lowest.

RECOMMENDATION: Click on individual teams to view their collaboration patterns in the charts on the right.

3 What is the impact of Copilot on various work patterns?

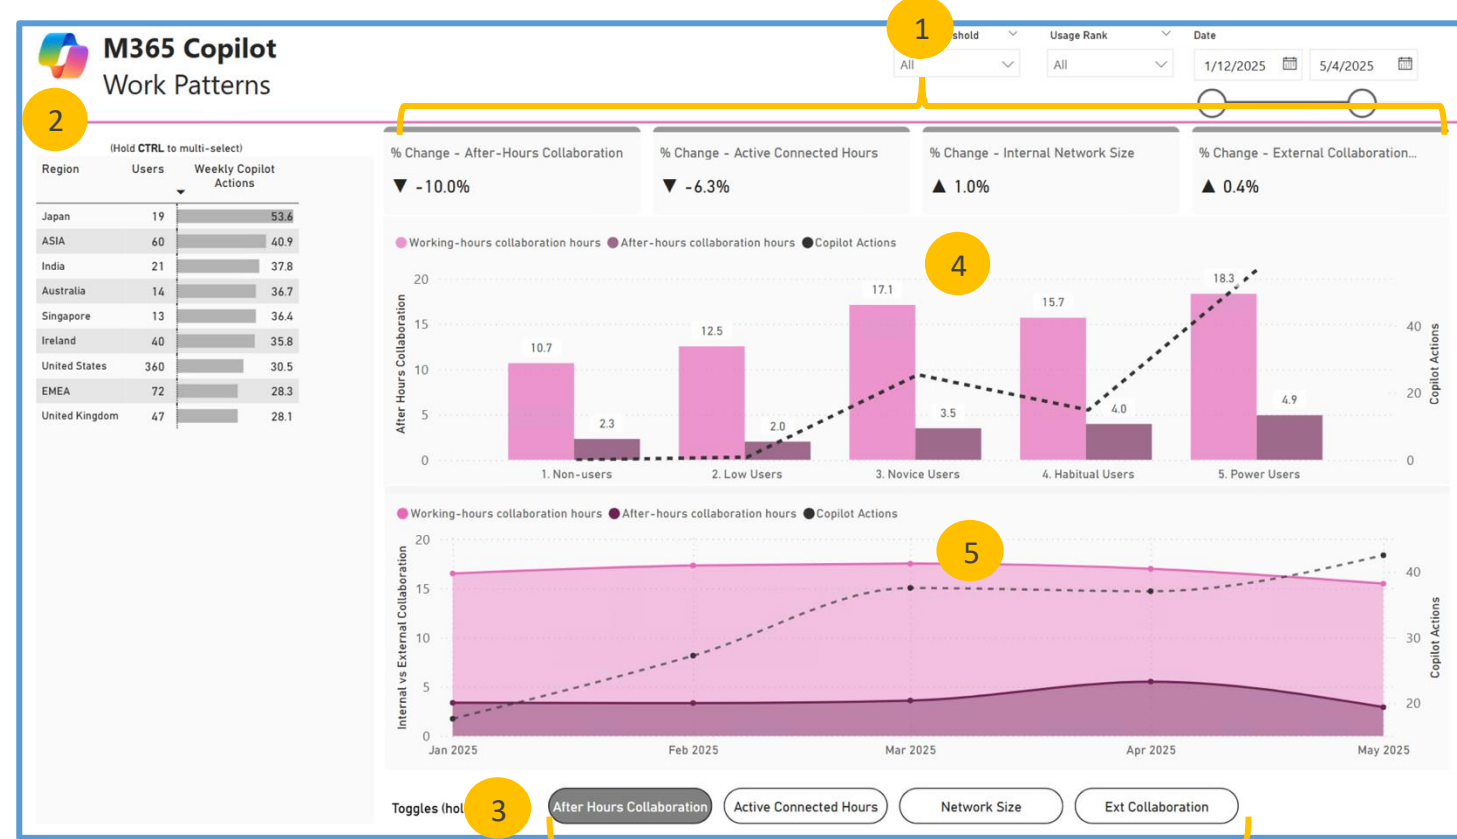
Use these buttons to toggle between metrics:

- After-hours collaboration
- Active connected hours
- Network size
- External collaboration

4 How does this metric vary across usage categories?

For this data, we see that power users have the highest working hours and after-hours collaboration.

INTERPRETATION: Is higher Copilot usage correlated with differences in collaboration patterns? Consider how increased Copilot usage may be impacting each individual metric. For example, do we see power users with higher external collaboration than other categories? Could this suggest a correlation between efficiencies gained from Copilot usage and an increased ability to serve external customers? Or, do we see a reduction in after-hours work among top users?



5 How do metrics change over time with Copilot adoption?

For this data, we see a slight decrease in working-hours collaboration as Copilot adoption increases over time, but no commensurate change in after-hours work.

INTERPRETATION: As with the top chart, consider the specific metric and the change you observe, along with Copilot usage. Is Copilot allowing users to spend time with customers and partners for roles that demand external interactions? Do we see a reduction in active connected hours as Copilot increases, suggesting Copilot may be positively impacting productivity and efficiency?

Note that impact on work patterns will vary by customer and roles. We frequently see patterns including after-hours work declining with increased Copilot usage and increased Copilot usage for those with higher collaboration hours, but your results may differ based on who is using Copilot and how they are using it.

RECOMMENDATION: Analyze patterns for power users over time - do they see the same pattern that we're observing for the overall population?

Mapping Features to Work Patterns

1

How has Copilot impacted meeting, email, and chat hours?

For this data, each of these metrics has declined over this reporting period.

Use the toggles in section 2 to switch between different metrics in the following charts.

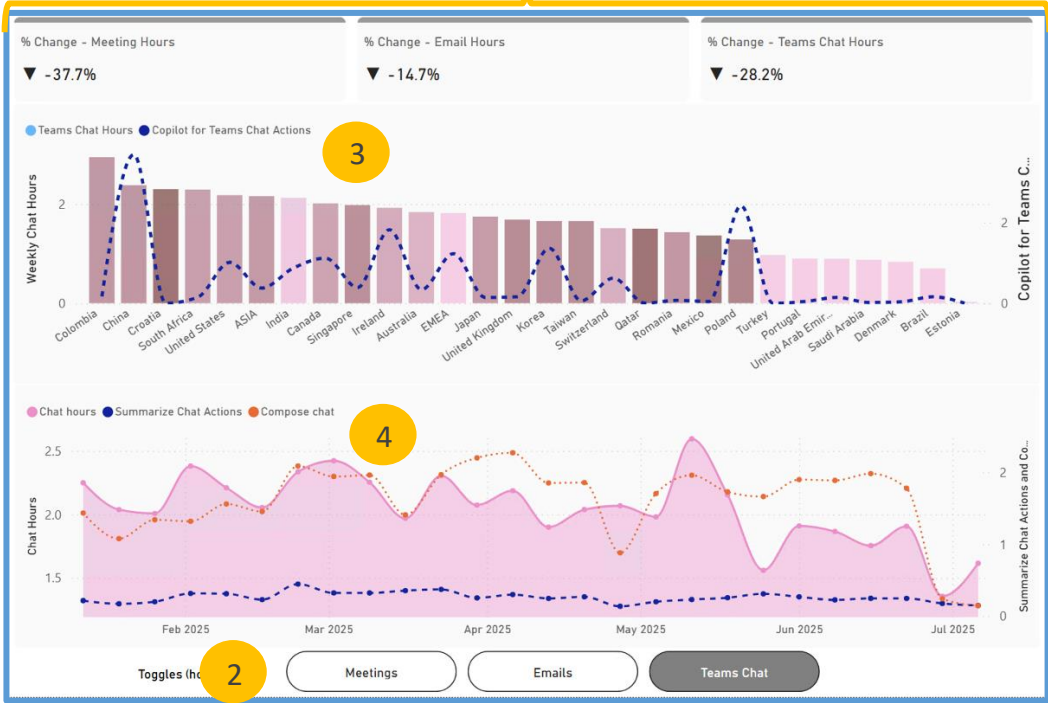
3

How has Copilot impacted the time spent in chat/email/meetings across teams?

For this data, we see China has the most Copilot chat actions on average, but also has one of the higher averages for weekly chat.

GUIDANCE: With all of these charts, it is important to remember that correlation does not equal causation, so results should be interpreted carefully. Similarly, sometimes new adoption of a Copilot feature can have an immediate impact on a particular behavior (like email or meetings), but it is important to keep observing the longer-term trend for a truer sense of impact. Note that impact on work patterns will vary by customer and roles. Your results may differ based on who is using Copilot and how they are using it.

RECOMMENDATION: Are there any teams where you see spikes or patterns that differ from other teams? Click on the bar chart for those teams to see their specific metrics. There may be an opportunity to learn more about the Copilot history for those teams – and always remember to look at group N size, as smaller groups often show greater volatility.



4

How has usage of specific Copilot features impacted chat/email/meeting hours over time?

For this data, we see that Compose Chat is used far more often than Summarize Chat, but there does not appear to be a strong correlation with chat hours.

INTERPRETATION: As with the top chart, consider the specific metric and the change you observe, along with Copilot usage.

RECOMMENDATION: Analyze patterns for power users over time - do they see the same pattern that we're observing for the overall population? In addition, if you see some features with little to no use, consider whether a targeted effort to publicize the value of those features would be beneficial.