2toLEAD

Telecommunications AI Agent Playbook

Crafting Conversations with Agents

WRITTEN BY | Kanwal Khipple & Richard Plantt

E-BOOK

TABLE OF CONTENTS

Introduction to Copilot Agents

What are Copilot Agents?

A Copilot Agent is an AI-powered digital assistant designed to help users' complete tasks, solve problems, and generate content through natural language interaction. Unlike traditional chat agents, Copilot Agents are conversational, goal-oriented, and context-aware, making them highly effective for a wide range of applications.

They leverage advanced language models to understand user intent, maintain coherent dialogue, and perform complex, multi-step tasks. These agents are often multimodal, capable of working with text, images, code, and more. These agents are typically integrated with tools, APIs, or internal systems to extend their functionality beyond simple conversation.

Real-world examples include Microsoft 365 Copilot, which assists with productivity tasks in Word, Excel, and Outlook; GitHub Copilot, which helps developers write and understand code; and custom-built Copilot Agents tailored for specific domains like customer support, education, or healthcare.

Building your own Copilot Agent allows you to tailor its capabilities to your unique workflows, integrate it with proprietary tools or data, and enhance productivity, creativity, and decision-making within your environment.

How Do Agents Differ from Traditional Chat Agents?

- **Smarter:** Uses advanced AI models (like GPT-4 or later) for deeper understanding and reasoning.
- **More Capable:** Can manage complex workflows, multi-step tasks, and dynamic user needs.
- **Context-Aware:** Maintains memory or session context to provide coherent, personalized assistance.

Why Build Your Own Agents?

- Tailor the agent to your unique workflows or business needs.
- Integrate with proprietary tools or data.
- Enhance productivity, creativity, and decision-making in your environment.

Privacy & Security

Copilot doesn't learn from individual user interactions in the way humans do. Instead, it relies on a vast amount of pre-existing data and continuous updates from its developers to improve its responses. This data includes books, articles, websites, and other text sources from trusted sources, ensuring the information provided is accurate and relevant.

While Copilot can remember the context of the current conversation to maintain coherence, it doesn't store personal data or learn from specific user interactions, ensuring your privacy and security. Additionally, users have the option to delete their Microsoft Copilot interaction history, which includes their prompts and the responses Copilot returns.

Integration with Organizational Data

Microsoft 365 Copilot's ability to leverage organizational data is a key differentiator. It integrates with various data sources within Microsoft 365, such as emails, Teams messages, SharePoint documents, and more. This integration allows Copilot to provide responses that are not only based on the user's prompt but also grounded in the data the user has access to within the organization. This ensures that the responses are relevant and tailored to the user's specific context.

For example, Copilot can:

- Summarize emails and documents from SharePoint and OneDrive.
- Retrieve information from Teams chats and meetings.
- Access calendar events to provide scheduling assistance.
- Utilize data from Dynamics 365 and other enterprise systems through Microsoft Graph connectors.

Enhancing Copilot with External Data

Copilot can also be extended to integrate with external data sources using Copilot connectors. These connectors allow organizations to bring in data from various systems, enhancing the AI-driven experience. For instance, you can use Microsoft Graph connectors to integrate data from non-Microsoft sources, providing a more comprehensive view and enabling more informed decision-making.

AI Use Cases for Corporate Communications

The Role of AI in Telecommunications

In the telecommunications industry, AI-powered agents integrated with Microsoft 365 are revolutionizing internal workflows and communications. Telecom companies have large, distributed workforces and fast-paced operations, ideal conditions for Copilot/AI agents to streamline planning, insights, and messaging.

Each "agent" leverages the rich ecosystem of Microsoft 365 (Teams, Outlook, SharePoint, Viva, Power Platform) combined with Copilot's AI capabilities to automate tedious tasks and provide intelligent assistance. The consistent card-like structure of these use cases is meant to help you scan and compare opportunities quickly. Built with low-code tools and Copilot Studio, your own business users and IT "makers" can develop and customize these solutions without extensive developer resources

Opportunities to Impact Corporate Communications

Corporate communicators have a powerful opportunity to transform their workflows with Copilot Agents by automating content creation, campaign planning, and internal messaging. These agents help streamline repetitive tasks like drafting announcements, summarizing insights, and managing multi-channel communications, freeing up time for more strategic work.

With AI-driven personalization, communicators can deliver more relevant, timely, and engaging content across platforms like SharePoint, Teams, and Viva. Ultimately, Copilot Agents enable communicators to shift from reactive messaging to proactive, data-informed storytelling that drives alignment and employee engagement.

Campaign Planning: Plan internal campaigns with AI-driven timelines, channel plans and KPIs for initiatives.

Audience Insights: Analyze employee sentiment from surveys and Viva Engage to tailor communication tone and timing.

Translation & Localization: Translate internal content into multiple languages with cultural nuance.

Crisis Comms: Draft urgent messages and coordinate multi-channel distribution with approval workflows.

Collect Feedback: Aggregate and summarize employee feedback, extracting themes and suggestion next step actions.

Leadership Voice: Help craft authentic messages by learning the style and tone of your leadership.

Communicator Roles Using AI

Communications Manager

Description: Oversees messaging within the organization, ensuring employees are informed, aligned, and engaged.

AI-Enhanced Productivity:

- Drafting and scheduling internal newsletters.
- Summarizing leadership updates or town halls.
- Analyzing employee sentiment from feedback channels.
- Automating responses to frequently asked questions in Teams or Viva Engage

Corporate Content Strategist

Description: Develops and manages content strategies that align with business goals and brand voice.

AI-Enhanced Productivity:

- Content calendar planning with AI-generated topic suggestions.
- SEO optimization for intranet articles.
- Repurposing long-form content into bite-sized formats.
- Tagging and categorizing content using metadata automation.

Crisis Communications Lead

Description: Manages communication during emergencies or sensitive situations.

AI-Enhanced Productivity:

- Drafting rapid-response messages with tone analysis.
- Monitoring sentiment and engagement in real-time.
- Creating scenario-based communication templates.
- Coordinating updates across multiple channels instantly.

Communication Analyst

Description: Tracks and reports on communication effectiveness and engagement metrics.

AI-Enhanced Productivity:

- Generating dashboards from Viva Insights and SharePoint analytics.
- Identifying trends in employee engagement.
- Creating executive summaries from raw data.
- Predicting communication impact using historical data

Internal Media Manager

Description: Manages the digital workplace hub, often built on SharePoint, Viva Connections or Engage.

AI-Enhanced Productivity:

- Auto-generating news posts from meeting transcripts.
- Curating personalized content feeds.
- Managing permissions and access requests.
- Suggesting layout improvements based on usage patterns.

Executive Communications Advisor

Description: Crafts messaging and presentations for senior leaders.

AI-Enhanced Productivity:

- Drafting speeches and talking points.
- Creating slide decks from meeting notes.
- Translating complex strategy into digestible narratives.
- Preparing Q&A briefs based on recent employee feedback.

Agent: Audience Insights Agent

Agent Name: Audience Insights Agent

Agent Description

Continuously analyze sentiment and feedback to help communications teams tailor their strategy. This agent digests data from sources like Viva Glint (employee engagement surveys), Viva Engage posts, or Microsoft Forms, and surfaces insights about how different employee segments feel.

It then recommends the optimal tone, timing, and targeting for upcoming communications. In a telecom company, for example, this agent can tell if frontline technicians are feeling disengaged about a new scheduling system rollout, or if call-center staff reacted positively to a recent policy change, enabling comms to adjust messaging accordingly.

Tools Used

Copilot Studio: For a conversational interface where a user can query insights.

Microsoft Viva Glint/Viva Engage: Use APIs to retrieve discussion threads.

Microsoft Forms/Excel: Used for collecting any ad-hoc survey data. The heavy analysis is done with AI: e.g., AI Builder's sentiment analysis and Azure OpenAI GPT for summarizing themes.

Power Automate: Used for flows connect to data sources (there are connectors for Viva Engage, Excel, and Glint data can be exported), then feed the data into the AI for processing.

Teams/SharePoint: Used to show the results can be shown in Teams chat or output as a SharePoint report.

Power BI (Optional): Used for visualizing sentiment trends over time.

Key Actions & Data Flow

The Audience Insights Agent performs sentiment analysis and audience segmentation on internal feedback, then translates that into communication advice:

Aggregate Multi-Channel Feedback: Pulls in textual feedback from various internal channels: e.g., the latest survey comments, posts and replies from Viva Engage groups (like an "All Employees" community or specific interest groups), or responses from a recent Microsoft Forms pulse survey. All this qualitative data (which could be thousands of comments) is aggregated for analysis.

Analyze Sentiment and Themes: Using AI, the agent analyzes the data to determine overall sentiment (positive, neutral, negative) and to extract the main themes or topics employees talk about. For example, it might find that "work-from-home policy" and "sales targets" are top discussion themes, with mostly negative sentiment around the former and mixed sentiment on the latter. The agent can categorize comments by emotion or intensity as well (identifying if people are excited, confused, frustrated, etc.).

Segment by Audience: If metadata is available (from Glint or Engage profiles), the agent can break down sentiment by department, region, or role. For instance, it may reveal that engineers are reacting more positively to a change than customer support reps, or that one office location has concerns not seen elsewhere. These insights help target messages more effectively to each group.

Recommend Tone and Timing: Based on the sentiment analysis, the agent provides recommendations for upcoming communications. For example, if the sentiment about a new network upgrade is apprehensive, it might advise to address the concerns candidly in the next announcement and use a reassuring tone. If employees seem overwhelmed by communications (e.g., engagement is dropping), it could suggest slowing down the cadence or sending messages at a different time of day when employees are more receptive.

How to Build the Agent

Connect Data Sources

Start in Power Automate. Build flows to collect data: e.g., a scheduled flow that runs weekly to fetch new Yammer messages (using the Viva Engage connector) and export Viva Glint survey comments. Store this data in a SharePoint list or Dataverse table. You can also input data on-demand via the agent (e.g., user pastes a chunk of text or uploads an Excel of survey results for instant analysis).

Sentiment & Theme Analysis

Use AI Builder sentiment analysis action to score each comment for sentiment (very positive to very negative). Additionally, use an OpenAI or Azure Cognitive Service for key phrase extraction or topic clustering. You might feed all comments into a prompt like "Summarize the top 5 themes mentioned and how people feel about each." This is exactly how Microsoft's internal team uses Copilot, asking it for thematic analysis of feedback.

Build the Q&A Interface

In Copilot Studio, create the Audience Insights bot. Give it topics like "Overall Sentiment" and "Segment Insights". When triggered, the bot will run the Power Automate flows to get the latest data and analysis. The dialog might allow the user to ask questions like "How did employees react to the new billing system announcement?" or simply "What's the mood this week?". Map these to actions that retrieve the prepared analysis or even run on-the-fly analysis.

Output Recommendations

Format the bot's answer to include not just data (e.g., "Overall sentiment 65% positive, 20% neutral, 15% negative"), but also an advice section. For example, if a theme of "confusion about new health benefits" is detected, the agent's answer could say: "Many employees are confused about the new health benefits (theme mentioned by 30% of responses). Consider sending a clarification email and an FAQ link.

Tone should be supportive and clarifying." These recommendation phrases can be generated by GPT based on rules or few-shot examples.

Test with Scenarios

Try queries in the Teams interface: ask the bot for insight after a major internal announcement and see if the advice makes sense. Adjust the prompt that generates recommendations to align with your company's style (you might include a sample "advice" output in the prompt for consistency). Ensure sensitive data is handled properly (the analysis likely should aggregate and anonymize comments).

Example Prompts

The Digital Communications Director at a telecom queries the Audience Insights Agent: *"How are employees feeling about our new remote-work policy?"*

The agent pulls the latest pulse survey comments and Viva Engage discussion on that topic. It reports: "Sentiment on the remote-work policy is 60% positive, 25% neutral, 15% negative."

Top themes:

1) Work-life balance improved (positive)
2) Concerns about team collaboration (negative)
3) Need for better home office support (neutral)."
4) The agent then advises: "Consider acknowledging collaboration concerns in your next update and highlight tools/training for remote teamwork. Overall tone should be appreciative of feedback and optimistic. Best to post the update early Monday when engagement is highest."

Armed with this, the Director can craft a message that directly addresses employee concerns and choose an optimal timing, likely resulting in better reception.

Over time, she notices that using these AI-generated insights leads to higher read rates and more likes on communications, confirming that aligning with the audience sentiment makes communications more effective.

Agent: Localization Agent

Agent Name: Multilingual Translation Agent

Agent Description

Automatically translate internal communications content into multiple languages and localize the tone for different regions or cultures. Telecommunications companies often operate in diverse regions (or employ staff from many linguistic backgrounds).

This agent ensures that important announcements, policy documents, or training materials reach all employees in their preferred language, with appropriate cultural phrasing. It saves enormous time compared to manual translation and guarantees consistency in messaging across languages. For example, if headquarters prepares an English memo about a new retail promotion, the agent can instantly provide the French, Spanish, or Japanese versions for regional offices, tailoring idioms or formality level as needed.

Tools Used

Copilot Studio: Used to implement the agent in Microsoft 365. The core is the Microsoft Translator service (accessible through the AI Builder prebuilt model or the Translator connector) which can translate text between dozens of languages.

Teams: Used as the agent interface as a message extension or

Copilot Agent (Optional): An alternate chat interface where a user submits text and selects target languages.

SharePoint: Serves as storage for original and translated documents (with a library column for language).

Power Apps (Optional): Can provide a simple form UI to input text and get translations. The solution remains low-code: using the drag-and-drop Translator action in a flow to perform the translation.

Azure OpenAI (Optional): This step can be added to adjust tone or terminology after translation for localization polish.

Key Actions & Data Flow

The Translation & Localization Agent makes translating content as easy as a conversation:

Multi-Language Translation: The user provides the content (plain text or an Office document) and specifies target languages. The agent can handle one or multiple languages in one go. By calling the Translator API, it produces quick translations.

For example, an HR policy update written in English can be translated into Spanish, French, and Hindi within seconds, all preserving the original formatting if it's a document. The AI Builder “Translate text” model detects the source language automatically and outputs the translated text5.

Cultural Localization: Beyond literal translation, the agent adapts the phrasing to suit local conventions. This might involve switching idioms, adjusting formality, or using region-specific terms. For instance, a casual English phrase like “Let’s get the ball rolling” might become the equivalent of “let’s get started” in another language if the idiom doesn’t translate.

Consistency and Glossaries: The agent can incorporate a company glossary or previously translated terms (e.g., product names or tagline translations stored in SharePoint) to maintain consistency.

For a telecom, that means if “Super Connect Router X1” is the product name, it stays in English or uses an approved localized name everywhere. The agent might query a list of known terms and ensure they are not translated or are replaced post-translation.

Output Formatting & Distribution: Finally, the agent delivers the translated content in the desired format. If the input was given as a Teams message or plain text, it replies with the translated text for each language labeled clearly. If it was a document (say a Word file), the agent could create new files for each language (like "Policy_Update_es.docx" for Spanish) and provide links. It might also post the localized announcement directly into regional Teams channels or send to the respective country managers via Outlook. All these can be automated steps.

Feedback Loop: The agent allows users (especially bilingual employees or local staff) to provide corrections if needed. It can learn from these edits by updating a glossary or translation memory. Over time, the translations become more tuned to the company's preferences.

By automating translation, companies ensure inclusivity, every employee can consume communications in their strongest language, which can improve understanding and engagement. It also enforces a unified message across the organization, as the translations are generated from the same source content instantly (no lag waiting for human translators).

How to Build the Agent

Create Translation Flow

In Copilot Studio, use the AI Builder: Translate text action inside a flow. For example, set up an instant flow that triggers when a user submits text (via a Power Apps form or a manual trigger in Copilot Studio). Use "Translate text into another language" and configure it for each target language needed. You can have multiple translate actions or loop through a list of languages. The AI Builder model automatically detects the source language and outputs the translation, which you can capture in variables.

Add Localization Refinements (Optional)

After each translation action, add an OpenAI (GPT) step with a prompt to refine the translation.

For instance: "You are a native French business writer. Improve the translation of the following announcement to sound natural and culturally appropriate for employees in France, while preserving the meaning." Feed it the raw French text and get the adjusted output. Do this for each language to localize tone.

Integrate Glossary (Optional)

If you have specific terms not to translate, you could implement a simple replacement step before translation (to mark those terms) or use the Translator custom dictionary feature if available.

Alternatively, after translation, search and replace any key terms using a stored mapping (Power Automate can do find/replace in text easily).

User Interface with Teams

Create a Teams message extension or a simple chat bot for the front-end.

For example, with Copilot Studio, you can make the agent respond to commands like “Translate” or even monitor when a message is posted in a particular channel and auto-translate it. A straightforward approach is a agents where the user says, “Translate this message.”

The bot then asks for the languages needed and the text (if not already provided), triggers the translation flow, and returns the results. Alternatively, a custom Adaptive Card in Teams could be used: the user enters text and checks boxes of target languages, and the card displays the translations upon submission.

Testing & Deployment

Test the agent with a variety of content, short announcements, long policy documents, technical jargon to see the quality. Verify formatting for example, line breaks and bullet points come through properly.

Ensure the translations are accurate by having native speakers review them initially. Once satisfied, deploy the agent for general use. You might pin it in a Teams channel (e.g., an “Internal Comms Toolkit” channel) or as an app in the Teams sidebar for easy access.

Example Usage

A Regional Communications Lead in a telecom company needs to send out a technical update about network upgrades to teams in both Brazil and Germany. She types the draft announcement in English into the Translation Agent (via a chat in Teams or an app interface) and selects Portuguese and German. The agent swiftly responds with:

Portuguese (Brazil): "Prezados colegas, ..." (followed by the full translated announcement, using local telecom terms in Portuguese),

German: "Liebe Kolleginnen und Kollegen, ..." (with formal greeting and context in German).

The agent's localization logic ensured that the Portuguese version uses a friendly tone common in Brazilian corporate communication, and the German one uses the proper level of formality.

It even adjusted a phrase: the English "We're excited to roll this out" became "Estamos animados para implementar isso" in Portuguese (a natural translation meaning "we are excited to implement this").

The communications lead reviews and with one click, tells the agent to distribute: it posts the Portuguese message to the Brazil office's Teams channel and emails the German version to the Munich office.

What used to take days coordinating with translation services is now done in a minute or two, entirely within their M365 environment.

Agent: Crisis Communications Agent

Agent Name: Rapid Response Agent

Agent Description

Provide an assistant for drafting and disseminating urgent communications during a crisis. In the telecommunications sector, crises can range from major network outages affecting customers, to internal emergencies or PR incidents. This agent helps communications teams and leadership quickly craft clear, concise messages and push them out via the right channels, following predefined workflows for approvals and escalation. The goal is to save precious minutes when time is critical, ensuring employees (and even customers, through internal comm relays) get timely, accurate information. It also maintains a consistent voice under pressure by using templates and AI guidance.

Tools Used

Microsoft Teams and Outlook: Used for dissemination, SharePoint for storing crisis plan templates.

Copilot Studio: Used for orchestrating the workflow (like approvals, multi-channel posting). It can be implemented as a Teams chatbot (for example, a communications manager can invoke it in a chat during an incident) or even integrated into an incident management system. Dialog to handle input (details of crisis) and guide the user through the process

Pre-approved message templates (e.g., an outage notification template, a data breach template) possibly stored in SharePoint or as prompt scenarios in Copilot Studio.

Approvals via Power Automate: before sending, the agent can route the draft to an executive for one-click approval (leveraging the Teams or Outlook approval adaptive card feature).

Distribution Connectors: Outlook (to send all-employee emails), Teams (to post in announcement channels), possibly SMS via a third-party connector if needed for critical alerts.

Key Actions & Data Flow

When a crisis hits, the agent will:

Gather Key Details: Through a quick Q&A, the agent asks the user for essential information: "What happened?" (e.g., "A fiber cable cut is causing an outage in downtown area"), "Who is impacted and how long?" (e.g., "500 enterprise customers, estimated 4 hours to fix"), "What's our immediate ask or guidance for employees?" (e.g., "Frontline teams should standby for updates; customer service should use workaround X"). This can be input via a form or chat prompts. Having this structured ensures nothing critical is omitted.

Draft the Message using Templates: Based on the incident type, the agent selects an appropriate template (for instance, an outage notification vs. a facilities emergency). It then uses AI to fill in the details provided into the template, adjusting tone as needed.

The draft message will typically include: a clear headline, the factual details, any action required by employees, and next steps or timing for updates. Because it's AI-assisted, the agent can refine the wording to be clear and calming, avoiding jargon or panic.

For example, it might produce: "Outage Update: A major fiber cut is currently disrupting connectivity for customers in . Our engineering teams are on it, and service is expected to be restored within 4 hours. What you need to do: Customer support should inform callers we're addressing the issue and offer to follow up. All other employees, please standby for another update by 3 PM.\" This saves the comms manager from writing from scratch under stress.

Approval Workflow: If required by policy, the agent routes this draft to the relevant leader (say, the VP of Communications or CTO) for review. Using Power Automate, it can send an Approval request card in Teams or an email with Approve/Reject buttons along with the message preview. This ensures oversight and that the tone is appropriate. Given time sensitivity, this can be configured to auto-send after a certain timeout if no response (or have a predefined list of approvers where any one can approve).

Multi-Channel Distribution: Once approved, the agent triggers notifications on multiple channels.

- It sends an urgent all-staff email via Outlook (maybe using a dynamic distribution list for all employees or a specific segment if the crisis is localized).
- It posts the message in a Teams "Announcements" channel or a company-wide Team, possibly tagging @All or specific groups (there's a "Company Communicator" app pattern that can be leveraged).
- If applicable, it can also integrate with SMS/phone alert systems (telecoms might have these for field engineers) via connectors or third-party APIs.
- It logs the communication in a SharePoint log or an internal incident management log for audit.

Escalation and Follow-ups: The agent doesn't stop at one message. It tracks if a follow-up is needed (e.g., if the resolution time passes or new info comes in). It can send reminders to the comms team: "It's time to issue an update – want me to draft it?" Similarly, if no resolution by a certain time, it might automatically escalate the issue update to a higher-level alert. These workflows ensure no crisis communication falls through the cracks.

Templates for Different Scenarios: The agent can handle multiple crisis types: safety/emergencies (like office closures during extreme weather), technical outages, PR issues, etc., each with slight variance in message style and channels.

By having this agent, telecom companies can respond faster and more uniformly during crises. The communications are sent out in a timely manner and are consistent in structure, which builds trust. Employees get the information they need to respond to customers or ensure their own safety without delay.

How to Build the Agent

Prepare Templates & Prompts

Compile the common crisis scenarios and draft message templates for each. Store these in an accessible form e.g., a SharePoint Online list or document library with columns for scenario type, template text.

Alternatively, encode these as system prompts in Copilot Studio: e.g., "If incident type = Outage, use this style... If Safety, use this style...". Having a library of content for the AI to pull from will make the drafts more accurate.

Setup the Chatbot (Copilot Studio)

Create a Crisis Comms agent. Define the dialog that triggers on keywords like "crisis" or a manual invocation.

The dialog will ask the user a series of questions (incident type, details, impact, etc.).

Use actions within the bot to store these answers and retrieve the appropriate template content.

Then use an action to merge details into the template and polish the text. Essentially, you prompt the AI: "Here is a template and here are the details, draft a clear urgent message." (A few-shot with examples of good messages could be included to guide tone).

Implement Approvals

Use the Approvals connector in Copilot Studio. After drafting the message, the flow creates an approval and posts it, for example, to Teams (or sends email) to the chosen approver.

The approver sees the draft message and clicks Approve or provides edits. If rejected or edited, you can loop back, perhaps the agent prompts "The approver suggested changes, please revise the details or message."

In Copilot Studio you can handle the approve/reject branches accordingly. If approved, proceed.

Multi-Channel Send

Connect to Outlook with the Send Email action addressed to all employees (could be a distribution list or all users in Azure AD).

Also, use a Teams connector (like "Post a message in a chat or channel") to publish the announcement in a Teams channel. If SMS integration is needed, either use a connector (Twilio, etc.) or rely on an existing system triggered by an email. Ensure these actions are setup in the flow after approval.

Testing in a Sandbox Environment

Simulate a few crisis scenarios (non-actual ones!) to verify the flow.

Example: a fake “building evacuation drill” scenario, feed it in, see that the agent drafts something sensible like an evacuation notice. Test the approval by having a teammate click the Approve button and confirm messages go out (perhaps to a test group).

Fine-tune the templates or prompts if the style isn’t quite right (e.g., too verbose or too terse). Also test edge cases: what if user says “no approval needed” agent should skip that step.

Once good, you can document the process and train the comms team on using the bot during real incidents.

Example Usage

It’s 7:00 AM, and a major network outage occurs unexpectedly, affecting customers and potentially a segment of employees. The telecom’s Internal Communications head uses the Crisis Comms Agent via their phone’s Teams app. They invoke the bot: “Network outage in Midwest”.

The agent asks for key information, and they quickly thumb in responses. Within a minute, the agent drafts an outage alert message. The Communications head hits approve (or maybe their director in charge does so promptly on Teams).

Immediately, an alert email goes out to all staff: “Urgent: Network Outage in Midwest Region: All customer-facing teams please be aware.” The message includes the core details and guidance.

Simultaneously, a Teams announcement is posted in the “Company Wide” channel with the same info. Field technicians and support reps get the news on their devices and are prepared to answer customer inquiries.

Because the agent had a template and filled it with the specifics, the message is accurate, calm, and actionable, avoiding any misinformation.

30 minutes later, the agent reminds to send an update. The issue is still ongoing, so the agent helps draft a follow-up with progress and an ETA for fix.

By the time the crisis is resolved, the communications were handled smoothly. Visualize the time saved: what used to take maybe 15-20 minutes of frantic writing and manual emailing now took under 5 minutes, potentially reducing customer impact by keeping everyone informed.

Agent: Leadership Voice

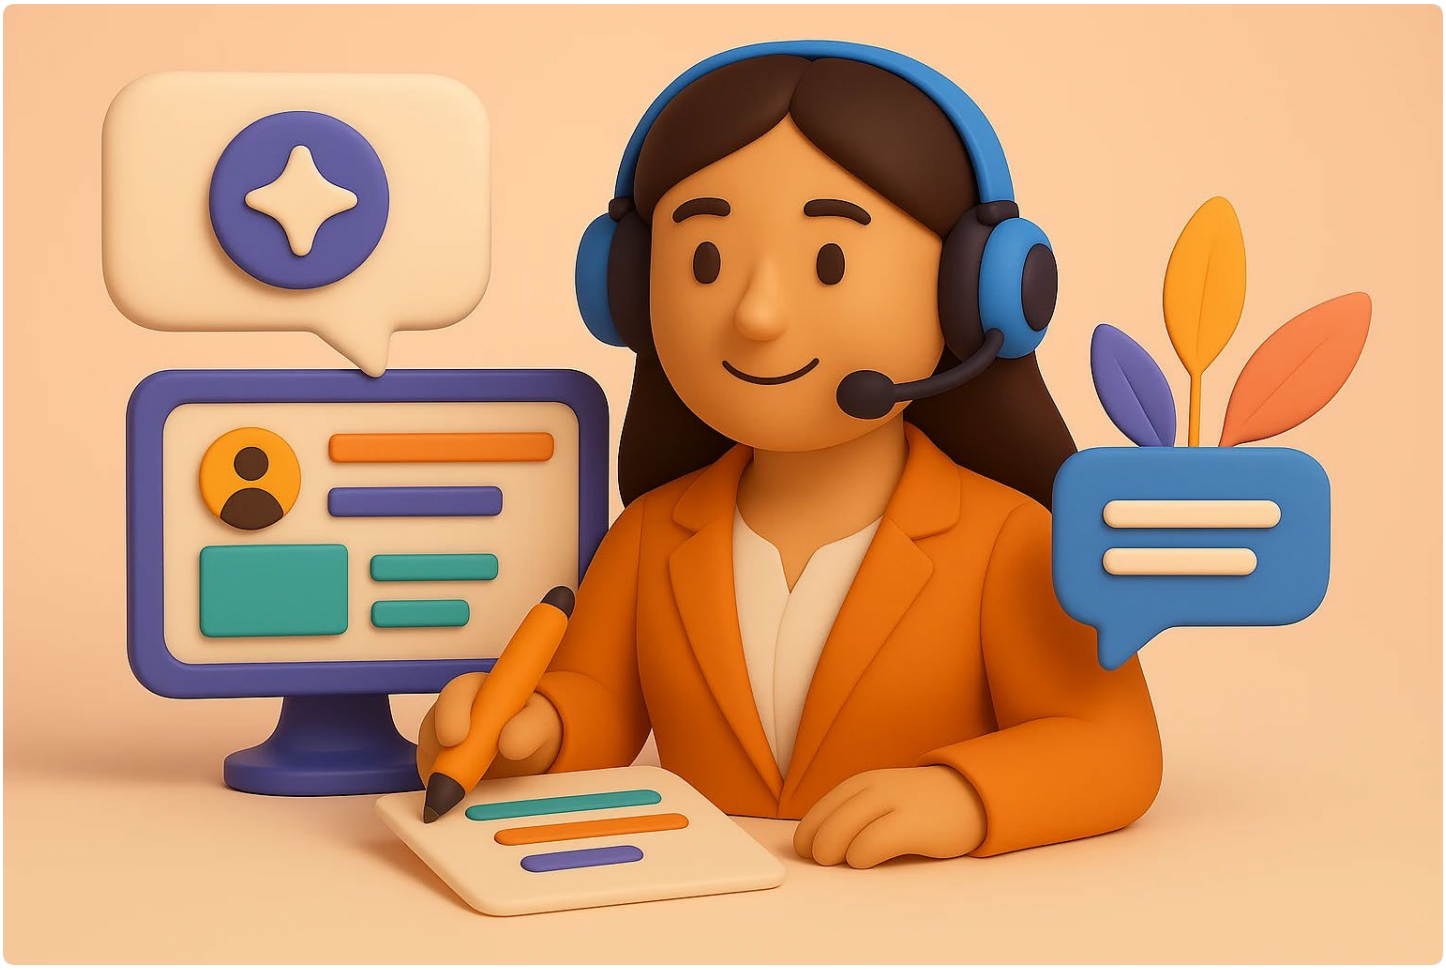

Agent Name: Executive Voice Agent

Agent Description

Assist executives in crafting authentic, personalized communications that reflect their unique voice and the company's values. Often, communications professionals ghostwrite emails or speeches for leaders (CEO messages, CTO tech updates, etc.), trying to mimic their style. This agent streamlines that process by learning from a leader's past speeches, emails, and posts to generate drafts that sound like the leader wrote them.

It ensures consistency with the leader's tone (be it inspiring, pragmatic, humorous, etc.) and weaves in key themes like company values or strategy points that the leader often emphasizes. In context, think of a CTO sending a quarterly update on network innovation. The agent can help include that CTO's favorite analogies or vision statements, so the message feels genuine and engaging to employees.

Tools Used

This agent heavily uses Microsoft 365 Copilot capabilities in Outlook, Word, and Teams. To build it, you would use Copilot Studio to create a custom conversational agent for the leader or comms team. Data sources for grounding the agent include:

SharePoint: Used to store the leader's prior communications: e.g., a collection of past company-wide emails, transcripts of town hall speeches, blog posts on the intranet, etc., (with the leader's permission). These serve as a knowledge base for the agent to pull style and content from. Other assets to help the agent:

- The company's mission/values statements and recent strategic narratives (which are then used to reinforce important messages).
- A style guide if available, describing the desired tone (though the actual examples are most useful).

Copilot Studio: Used as the foundation for the agent. The agent retrieves relevant snippets from past content (for context or even to directly quote/repurpose if appropriate) and then uses GPT to compose the new message. All of this can be done with low-code configuration in Copilot Studio by defining a plugin or data connection to those documents and letting the AI ground on them.

Key Actions & Data Flow

The Leadership Voice Agent's flow is about input -> draft -> refine, capturing the leader's style.

Input Understanding: The user (maybe the leader themself or their communications aide) tells the agent what they need to write. For example: "Draft this month's CEO newsletter about our Q3 results and thank the teams for their hard work, also mention the 5G rollout success and emphasis on customer focus."

The agent will parse this request, possibly asking follow-up questions if details are missing ("Any specific customer story to include?").

Retrieve Leader's Past Highlights: The agent searches the curated knowledge base of the leader's prior communications for anything relevant to the topic. It might find, for instance, that the CEO spoke about customer focus in last quarter's email and always uses the phrase "putting customers at the center of everything we do," or that the 5G rollout was mentioned in an earlier strategy memo with some impressive statistics. The agent pulls those reference points (this is done via the Copilot retrieval plugin hooking into SharePoint or email archives). This ensures factual grounding and stylistic clues.

Generate Draft in Leader's Voice: Using the retrieved content and the prompt, the agent composes a draft of the message. It will emulate the writing style, if the leader often opens with a personal anecdote, the draft might start with one (perhaps created by the AI unless a real one is specified).

If the leader prefers bullet points for achievements, it will do that. Essentially, it uses the leader's own words from the knowledge base to steer the tone. Microsoft's Copilot in Word already allows users to "rewrite this to match the style of my previous communications", and this agent automates that from the start.

One could even instruct it: "Use a confident, optimistic tone, as CEO usually does, and keep it around 300 words." The draft would incorporate those elements.

Iterative Refinement: The agent can accept feedback. The user might say, "Make it shorter and punchier," or "Add a bit of humor like CEO did in the January note." The AI will adjust the draft accordingly, perhaps by removing some corporate buzzwords and adding a light joke if appropriate. This iterative loop continues until the user is happy. The agent might also check against company guidelines (e.g., ensuring no confidential info is included inadvertently and that the tone aligns with brand voice guidelines).

Final Touches, Polishing and Format: The agent suggests a subject line (for an email) that resonates, maybe even A/B options as Copilot did for subject lines internally. It proofreads and ensures clarity.

Because it's aware of the leader's typical phrasing, it might catch if something "doesn't sound like them" and propose an alternative. Essentially, it serves as both a writer and an editor tuned to one person's style.

When done, the leader or comms team has a ready-to-send message that feels authentic and personal. This saves time (the agent did the heavy first draft) and serves as a writing coach. Over time, the leader might even use it to refine their style by seeing suggestions.

How to Build the Agent

Collect Leader's Writing Samples

Work with the executive's team to gather a set of at least 10-20 substantial pieces of their communication. This could be past emailed memos, internal blog posts, speech transcripts, etc.

Make sure to get approval, as this can be sensitive data. Store these in a dedicated SharePoint folder or a OneNote that Copilot can access.

Alternatively, use a Copilot plugin to connect to the mailbox (with appropriate scope) to fetch these on the fly. The key is you want the AI to have access to these texts to learn the style and to possibly reuse key phrases.

Define the Agent in Copilot Studio

Create a new Leadership Voice Agent and specify the persona (e.g., an assistant that is an expert mimic of the leader's tone and knowledgeable about company strategy). Use the grounding feature to connect the SharePoint folder of past communications as a data source. This way, the agent can cite or use info from those documents during conversations.

Dialogue & Prompts

In the dialog manager, create a structure for requests like “Draft an email” or “Help with a speech”.

You might allow the user to select what type of content (email, Teams post, speech, etc.) and the key points they want included. Then the agent’s logic node would assemble a prompt for the LLM that includes relevant retrieved text. A possible system prompt: *“You are a writing assistant for CEO. Write in CEO’s style: refer to themselves as 'I', use a motivational and appreciative tone. Include references to company values when relevant. Here are some of their past quotes for style reference: \[insert snippets]. Now draft a message about {topic with points provided}.”* This mixes retrieval with instructions.

Leverage Copilot Features (Optional)

You can integrate Microsoft Editor checks or Copilot refinements: for example, after the AI generates a draft, run it through the Copilot in Word API (if accessible) to optimize clarity or check alignment with writing style guidelines. However, the model itself should suffice for style mimicry if given enough examples.

Testing with the Comms Team

Try out drafting a known past message to see if it comes close to the original (a fun exercise to validate style capturing). Also test with new hypothetical topics.

Pay attention to whether it’s accurately reflecting the leader’s voice, does it use phrases the leader often uses? Does it avoid things the leader wouldn’t say? You might need to add a few “do’s and don’ts” in the prompt (like “CEO often says 'team', never uses 'employees', prefer that term” etc.).

Ensure factual accuracy by seeing that if it pulls data (like numbers or achievements), it’s from the provided sources or updated info given by the user, not hallucinated.

Deploy

The final usage could be the comms person typing a request in a Teams chat with the agent, or directly within Outlook. If using Outlook, you could deploy the agent such that it appears in the Outlook Copilot pane for that leader’s account.

Alternatively, the comms team member uses the agent in Teams, then copies the draft into an email for the leader to send. Both are viable. Provide training or documentation to the executive’s team on how to prompt the agent effectively (e.g., “tell it the main message and any specific tone cues”).

Example Usage

The CEO of the telecom wants to send a message to all employees about the successful completion of a major project (say, a nationwide network upgrade) and inspire them for the next quarter. In the past, the CEO's communications often highlight personal gratitude and tie achievements back to the company's mission of connecting people.

The communications director uses the Leadership Voice Agent: *"Draft a 300-word email from Jane Doe (CEO) celebrating that we finished the 5G expansion project ahead of schedule, thanking the network and operations teams, and connecting it to our mission of innovation and customer commitment. Include a forward-looking note about upcoming challenges."*

The agent retrieves a line from a prior speech where Jane Doe said, *"Our mission is to connect people in ways never thought possible,"* and a phrase from a previous email, *"I couldn't be prouder of the teamwork we've shown."* It then produces a draft email along the lines of:

"Hello Team,

When I reflect on our journey this quarter, one word comes to mind: proud. We've not only expanded 5G across the country ahead of schedule, but we did it while upholding our core mission – connecting people in ways never thought possible. I want to personally thank each member of our Network and Operations teams for their dedication and ingenuity. …"*

It continues with the appreciative tone, perhaps mentions "as CEO, I believe this kind of achievement exemplifies our value of customer commitment," and closes with an optimistic look to the future challenge, using wording the CEO often uses like "onward and upward."

The comms director sees the draft and is amazed that it "sounds just like something Jane would say."

They might ask the agent to *"add a light-hearted note about celebrating this win,"* and it adds a line about taking a moment to celebrate with your teams today. After a quick review, the draft is sent to the CEO, who finds minimal changes needed.

In fact, the CEO notes the message "captures my thoughts exactly and even reminded me to mention our mission, great!" The open rates on this email are high, and employees remark that the CEO's notes always feel genuine. In one instance, using a tone-matching AI like this led to a ~30% increase in engagement with leadership communications3, as people resonate with the authentic voice.