

BENCHMARKING zoom AI

Productivity and Value Benchmark vs. Standalone and Suite-Integrated AI Assistants

(ChatGPT, Microsoft Copilot, Claude, Google Gemini)



Lead Author **Kevin Kieller**, enableUC

enableUC

EXECUTIVE SUMMARY

This report benchmarks Zoom AI against ChatGPT, Microsoft Copilot, Claude, and Google Gemini across ten practical knowledge-work scenarios, including meeting summaries, action-item extraction, grounded Q&A, document summarization, project-status updates, and cost/value calculations. The evaluation was not designed as a pure model comparison. Instead, it measured four enterprise-relevant factors: response relevance, time to insight, user experience, and cost. Across the test set, all tools were generally capable, but performance varied significantly when tasks required fast access to work context spread across meetings, documents, chats, and email.


On response quality, Zoom achieved the highest average accuracy score at 4.7 out of 5, ahead of ChatGPT, Gemini, and Copilot at 4.6, and Claude at 4.5. For completeness, Zoom, Claude, and Copilot tied for the highest average score at 4.5, followed by ChatGPT at 4.4 and Gemini at 4.3. These results suggest that overall answer quality was strong across the tools, with relatively narrow differences in accuracy and completeness.

The largest performance gap appeared in time to insight. Zoom completed all ten test cases in a total measured time of 100.2 seconds, compared with 181.05 seconds for Gemini, 245.3 seconds for ChatGPT, 330.6 seconds for Claude, and 364.8 seconds for Copilot. In these scenarios, Zoom's direct access to required work context, without manual document retrieval, made it roughly 3x to nearly 5x faster than the other tools. This matters because the benchmark included real workflow overhead: finding context, switching tools, retrieving files, and assembling source material. In day-to-day enterprise use, that overhead often determines whether an AI assistant delivers practical productivity gains.

The report concludes that Zoom's strongest differentiator is not simply model quality, but its ability to consolidate and use relevant work context through first-party capabilities and built-in connectors. This includes content from Zoom meetings, third-party meetings, in-person sessions captured through Zoom's My Notes, documents, chats, and email. That context advantage enabled fast, relevant responses while supporting a strong user experience.

The findings suggest that organizations may be able to realize meaningful productivity value from Zoom without adding separate, overlapping AI tools.

For enterprise buyers, the key takeaway is clear: in collaboration-centric work, context access, workflow integration, and time to insight can matter as much as raw model capability.



The report concludes that Zoom's strongest differentiator is not simply model quality, but its ability to consolidate and use relevant work context through first-party capabilities and built-in connectors.

DOCUMENT OVERVIEW

This whitepaper presents a benchmark-style assessment of Zoom AI (previously called AI Companion) focused on knowledge-work scenarios, including meeting insight extraction, work-grounded Q&A, and document-oriented tasks. The objective is not to crown a single "best" AI model in the abstract, but to evaluate practical productivity outcomes specifically where answers must be grounded in collaboration artifacts such as meeting content and organizational documents.

AI assistant adoption is accelerating, and technology buyers face a crowded landscape with overlapping claims and limited "apples-to-apples" comparisons. To support decision-making in a complex marketplace, this report details hands-on results organized across four evaluation pillars:

1 Response relevance and contextual grounding	2 Time to insight including setup overhead	3 User experience and workflow integration	4 Cost and risk of duplicative spend
---	--	--	--

A prompt library is included to help readers replicate key scenarios in their own environments, recognizing that results can vary by configuration, permissions, and tool versions.



TABLE OF CONTENTS

- Executive Summary 2
- Document Overview 3
- Table of Contents 4
- 1. Introduction 5
 - 1.1 How to Read This Report 5
 - 1.2 Market Context: The Rise of AI Assistants for Knowledge Work 5
 - 1.3 The Problem: Disconnected Tools and Productivity Leakage 5
 - 1.4 Thesis: In-Workflow, Context-Aware AI Can Outperform Standalone Tools for Collaboration-Centric Work 5
 - 1.5 Intended Audience, Scope, and Out-of-scope Elements 5
- 2. Solution Overview: Zoom AI 6
 - 2.1 What Is Zoom AI 6
 - 2.2 Core Capabilities and High-Value Use Cases Evaluated 6
- 3. Competitive Landscape and Comparison Approach 7
 - 3.1 Tools Included in This Evaluation 7
 - 3.2 Evaluation Pillars and Scoring Model 7
- 4. Methodology 8
 - 4.1 Test Environment and Setup 8
 - 4.2 Test Case Design 8
 - 4.3 Metrics, Data Capture, and Analysis 8
- 5. Findings: Quantitative Results 9
 - 5.1 Response Relevance and Quality Scores 9
 - Accuracy 9
 - Completeness 10
 - 5.2 Time-to-Insight Results 12
 - 5.3 Cost Comparison 13
- 6. Findings: Qualitative Insights 14
 - 6.1 Workflow Integration and User Experience 14
- 7. Comparative Analysis 16
 - 7.1 ZoomMate 16
 - 7.2 vs. ChatGPT (Standalone AI) 16
 - 7.3 vs. Claude (Standalone AI) 16
 - 7.4 vs. Copilot (Suite-Integrated AI) 17
 - 7.5 vs. Gemini (Standalone or Suite-integrated AI) 17
- 8. Conclusion 18
- About This Publication 19
- Disclosure, Disclaimers, and Definitions 19
- Appendices 19
 - Appendix A: Assumptions and Limitations 19
 - Appendix B: Detailed Test Cases (Prompt Library) 20
 - Appendix C: References 21
 - Appendix D: About the Author 21



1. INTRODUCTION

1.1 How to Read This Report

This section provides market context, current AI assistant challenges, and the overall thesis that drove this report.

Sections 2 through 4 define what was evaluated and describe the test design, including the scenarios, metrics, and controls used to support fair comparisons.

Sections 5 through 7 summarize the results and observed differences across tools, highlighting both strengths and trade-offs. Finally, Section 8 presents overall conclusions.

1.2 Market Context: The Rise of AI Assistants for Knowledge Work

AI assistants are increasingly used for common knowledge-work activities such as summarization, drafting, Q&A, and synthesis across meetings, documents, and messages.

Effectively evaluating the value of enterprise AI assistants requires criteria extending beyond model output quality alone. This is especially true given many employees struggle with multiple AI assistants, each often providing only a partial, incomplete solution.

Organizations should consider workflow fit, how easily an assistant can access the right context, governance and security requirements, and the total cost of adoption, including licensing and administrative overhead.

1.3 The Problem: Disconnected Tools and Productivity Leakage

Many knowledge workers run meetings in one system, store files in another, and turn to AI in a separate, standalone interface.

This fragmentation can introduce friction, requiring manual copy/paste, ad hoc context uploads, inconsistent outputs, and uncertainty about where data goes and how it is handled.

These overheads can increase “time to insight,” reduce trust, and slow adoption, particularly for meeting-centric work where context is distributed across agendas, transcripts, chat, follow-up emails, related documents, and often more than one meeting platform.

1.4 Thesis: In-Workflow, Context-Aware AI Can Outperform Standalone Tools for Collaboration-Centric Work

ZoomMate, Zoom's newly launched AI teammate, is positioned as an in-workflow assistant intended to reduce context-switching by grounding responses in Zoom meetings, documents, emails and drives, with first-party support for common external repositories such as Google Drive, Gmail, OneDrive, Outlook/Exchange, and Slack. ZoomMate also provides connectors for CRM, HR, and ticketing systems; however, these line-of-business application connectors are beyond the scope of this paper.

Further, My Notes by Zoom allows meeting transcripts and summaries to be captured and consolidated even for meetings held on other platforms (Webex, Google Meet, Microsoft Teams) and in-person.

The central hypothesis explored in this paper is that better access to relevant work context, combined with tighter workflow integration, can improve practical outcomes such as response relevance, task completion time, and overall user experience when compared to standalone AI assistants. For suite-integrated AI assistants, context, time to insight, and user experience are key factors that differentiate value.

1.5 Intended Audience, Scope, and Out-of-scope Elements

This report is intended for enterprise IT and collaboration leaders, procurement teams, security stakeholders, and business decision-makers evaluating AI assistants for productivity.

The scope is intentionally practical: it focuses on AI assistance for meeting- and content-centric work, and compares tools across relevance, time to insight, user experience, and cost considerations. It is not a deep model-benchmarking paper (e.g., standardized academic test suites), nor is it designed to evaluate coding performance or consumer use cases that fall outside typical business requirements.

2. SOLUTION OVERVIEW: ZOOM AI

2.1 What Is Zoom AI

Zoom AI is Zoom's umbrella of AI products and features, including the newly launched ZoomMate and Zoom's AI note-taker, My Notes, along with the AI features built into Zoom Workplace like meeting summaries, meeting questions, chat compose, and many others. They are designed to help users obtain answers and generate outputs grounded in their collaboration context, including meetings, chats, emails, and documents.

Zoom's My Notes is a unique personal note-taking tool that captures meeting context across any platform, whether it's a Zoom meeting, an in-person session, or a meeting on a third-party service like Microsoft Teams. By centralizing these insights regardless of where they happen, My Notes provides Zoom AI with richer, more comprehensive context to quickly generate better summaries and more accurate action items.

Zoom AI capabilities are available via a web interface and are integrated into the Zoom Workplace application. Zoom AI features are included at no additional cost with paid Zoom Workplace plans.

2.2 Core Capabilities and High-Value Use Cases Evaluated

The scenarios evaluated in this report emphasize common, high-frequency knowledge-work needs. These include extracting meeting insights (summaries, decisions, and action items), producing follow-up notes, and synthesizing themes across multiple meetings. The evaluation also includes document-centric tasks such as summarizing and refining content, drafting structured outputs (e.g., executive briefs), and answering questions grounded in meeting transcripts, summaries, and stored content. Throughout the tests, a key lens is "time to insight", how quickly a user can move from a question to an output that is usable in the next step of their workflow.



Organizations should consider workflow fit, how easily an assistant can access the right context, governance and security requirements, and the total cost of adoption, including licensing and administrative overhead.

3. COMPETITIVE LANDSCAPE AND COMPARISON APPROACH

3.1 Tools Included in This Evaluation

This evaluation compares Zoom's AI capabilities with a set of widely used AI assistants representing both standalone tools and suite-integrated approaches. The goal is to reflect options buyers are actively considering, while acknowledging that exact capabilities vary by licensing tier, tenant configuration, and release cadence.

- Zoom AI
- OpenAI ChatGPT (standalone)
- Microsoft Copilot (suite-integrated)
- Anthropic Claude (standalone)
- Google Gemini (suite-integrated / standalone as applicable)

3.2 Evaluation Pillars and Scoring Model

Findings are organized across four primary pillars.

1

Response relevance

assesses accuracy, completeness, and how well an answer is grounded in the specific work artifacts associated with the prompt.

2

Time to insight

measures the elapsed time from prompt to a usable output, including setup overhead such as locating source material, switching tools, and re-prompting.

3


User experience

considers workflow integration factors. How easily users can access context, iterate on outputs, and (where available) verify sources through citations. The evaluation of this pillar is the most subjective.

4

Cost

considers incremental licensing and administrative overhead, including the risk of duplicative spend when organizations adopt multiple assistants with overlapping capabilities.



4. METHODOLOGY

4.1 Test Environment and Setup

The employed methodology was designed to reflect how enterprise users use AI assistants in day-to-day work, rather than in isolated lab conditions. A representative work corpus was defined to support grounded scenarios (i.e. meeting transcripts or summaries, documents, emails, and content stored in shared drives), and consistent prompts were used across tools.

All tests were conducted on a "typical" corporate laptop (Lenovo, ThinkPad X1 Carbon Gen 10 with 32 MB RAM, running Windows 11) and connectivity via WiFi 6 with a 1 Gbps internet connection.

Accuracy and completeness scores were assigned by panel participants without knowing which tool produced which output.

Timing from prompt to response was measured using a consistent approach. In the final presented results, a standard 15 seconds per document retrieval was ascribed to tools that did not have direct access to the required context. This approach was designed to reflect an efficient copy-and-paste operation, removing user variability (some users took significantly longer to find and retrieve specific documents). Further, as even AI assistant tools with connectors often required more time to retrieve source context from external platforms, we did not want to unfairly burden specific tools.

4.2 Test Case Design

Test cases were designed to be balanced across meeting-centric and document-centric tasks, reflecting common enterprise workflows.

The prompt set includes direct fact lookup, summarization, synthesis across multiple sources, action extraction, drafting, and "what changed?" comparisons. (*See Appendix B*)

A scoring rubric was applied to support consistent evaluation across tools, while still allowing for qualitative notes where a strict numeric score would be insufficient to capture user experience differences. Numeric scores are presented in the quantitative analysis section, and a summary of user experience notes is captured in the qualitative insights section.

4.3 Metrics, Data Capture, and Analysis

Both quantitative and qualitative measures are used.

Quantitative metrics include time-to-first-usable answer, response relevance scored using a 1–5 scale across accuracy and completeness, and costs.

Where tools provided a selection of multiple models, tests were run multiple times to avoid penalizing a "fast" result for relevance, or needlessly selecting a "thinking" model and incurring a time penalty. We intended to optimize model selection, even if individual users may not do this.

Qualitative assessment captured perceived workflow friction and overall user experience, based on notes captured during execution of the tests.






5. FINDINGS: QUANTITATIVE RESULTS

5.1 Response Relevance and Quality Scores

Response relevance and quality are reported below by tool and scenario type, focusing on accuracy, completeness, and contextual fit for work-grounded prompts.

ACCURACY

The score for accuracy focused primarily on whether any incorrect or “hallucinated” facts were included in the AI-generated outputs.




All the evaluated AI tools scored well in terms of accuracy, with Zoom AI achieving a slightly higher average score, based on results from the judging panel.

Test Case	Focus	Zoom	ChatGPT	Gemini	Claude	Copilot	Notes
TC-01	Meeting summary	4	4	4	4	4	Very unstructured discussion
TC-02	Action items	4.75	4.25	4.3	4.5	4.25	ChatGPT and Copilot most confident (overly?); Claude least confident
TC-03	Decisions	4.5	4.5	4.5	4.5	4	Copilot had terse decisions; Zoom itemized
TC-04	Meeting question	5	4.5	5	5	5	ChatGPT lacking who
TC-05	Cross-meeting	5	5	5	5	5	Accurate summaries
TC-06	Email follow-up	5	5	5	4	5	Claude (Haiku only) provided overly aggressive due dates
TC-07	Doc summary	4.8	4.2	4.7	4	4.9	Claude and ChatGPT missed "not specified"
TC-08	Policy process	4.5	4.5	4.5	4.5	4.5	All generally accurate
TC-09	Project status	4.125	5	4.25	4	4.75	Claude exceeded word count
TC-10	Math	5	5	5	5	4.9	Copilot showed incorrect number in calcs but did not impact results
Avg Score		4.7	4.6	4.6	4.5	4.6	

With respect to the meeting summary evaluation (TC-01), the meeting discussion, as captured in the transcript, was very unstructured, and, as such, evaluators had difficulty assessing what "high" accuracy looked like; this resulted in lower scores across the board.

COMPLETENESS

Completeness scores focused on whether all key points, decisions, action items, or other relevant points were included in the AI-generated results.



Zoom AI Companion, Claude, and Copilot all achieved the highest average completeness scores. ChatGPT and Gemini posted lower scores, but all achieved acceptable scores.


Test Case	Focus	Zoom	ChatGPT	Gemini	Claude	Copilot	Notes
TC-01	Meeting summary	4	4	4	4	4	All had different highlights
TC-02	Action items	4.5	4.5	4	4.5	4.5	Gemini had less items
TC-03	Decisions	5	4.25	3.5	4	3.5	Zoom separated decisions
TC-04	Meeting question	4.75	4.5	4.75	4.8	4.5	ChatGPT and Copilot lacked detail in summary
TC-05	Cross-meeting	5	5	5	5	5	Stylistic differences
TC-06	Email follow-up	5	4.5	5	4.5	5	ChatGPT generic mtg name; Zoom assigned Sarah to coordinate; Claude wordy; Copilot offered to create Outlook email
TC-07	Doc summary	4.5	4	4.7	3.8	4.8	Claude and ChatGPT missed "not specified"
TC-08	Policy process	4	4	4	4	4	Limited consensus on what "good" looked like
TC-09	Project status	4	4.5	4	5	4.75	Brevity vs detail
TC-10	Math	4.5	4.5	4.5	5	4.5	Claude created an interactive dashboard
Avg Score		4.5	4.4	4.3	4.5	4.5	

While results were on average acceptable, there were instances where the results omitted a detail that the judging panel felt was important. Reviewing and verifying output from any AI assistant is still a necessary best practice.



5.2 Time-to-Insight Results

Time to insight captures how long it took each AI Assistant to provide a usable answer after the prompt and required context was made available.



Zoom AI was significantly faster in generating usable output. In addition, because Zoom had direct access to the required context (meeting transcripts, documents, emails) regardless of platform (via built-in connectors and Zoom My Notes), the total “time to insight” for Zoom was 3x to 5x faster.

Test Case	Focus	Zoom	ChatGPT	Gemini	Claude	Copilot	Notes
TC-01	Meeting summary	11	80	74	45.3	95	Includes document retrieval time in totals
TC-02	Action items	8.9	11.9	7.8	56	22.8	ChatGPT connector does not access transcripts; Claude connector does; Copilot connector not available in our tenant
TC-03	Decisions	13	23.6	13.4	28.3	41.6	Meeting transcript
TC-04	Meeting question	7.2	7.5	9.6	26.1	8.8	Meeting transcript
TC-05	Cross-meeting	13.4	17	11.4	35.6	22.5	3 meeting transcripts
TC-06	Email follow-up	7	11.6	10.2	27.4	24	Meeting transcript
TC-07	Doc summary	10.5	32	11.4	11.2	28.6	Software spec pdf
TC-08	Policy process	11.2	31.6	13	24.2	59.4	Word doc
TC-09	Project status	12	12.1	14.15	50.5	27.1	Gantt chart + meeting transcript
TC-10	Math	6	18	16.1	26	35	No documents
Total Time		100.2	245.3	181.05	330.6	364.8	(11 documents)
Doc Retrieval		0	165	165	165	165	
Max Total		100.2	410.3	346.05	495.6	529.8	

5.3 Cost Comparison


Zoom includes many of its AI features at no additional cost with any paid Zoom Workplace plan. Zoom recently launched **ZoomMate** which provides additional agentic search, orchestration, and content creation features, available at additional cost for AI credits. We used the included AI features for this report.

Similarly, Google now includes its Gemini AI assistant with paid Gemini Workplace plans. Gemini is also available as a paid standalone AI tool (the AI Pro plan) at \$19.99 per user per month, which was the version used for this report.

Microsoft charges \$30 per user per month for Copilot as an add-on to any Microsoft 365 plan.

OpenAI's ChatGPT is available for \$20 per user per month, for the Plus plan, which was used in this report.

Anthropic's Claude is available for \$20 per user per month, for the Pro plan, which was used in this report.



**Zoom offers advanced AI capabilities through ZoomMate at an additional cost; however, the built-in Zoom AI features were used in this evaluation.*

While Microsoft, OpenAI, Anthropic, and Google do offer free versions of their respective AI assistants, these free tiers have significant limitations and are not positioned for business use in the types of benchmarks investigated in this report.



6. FINDINGS: QUALITATIVE INSIGHTS

6.1 Workflow Integration and User Experience

While all leading AI Assistants leverage extremely capable LLMs (large language models), these LLMs can answer questions only based on their extremely extensive, but generalized, training data.

In a business setting, the most valuable AI assistant is the one that can access specific organizational context. That context often lives in conversations, especially formal meetings, as well as pre-meeting, in-meeting, and post-meeting chats and documents.


The challenge with context is understanding what source materials your AI assistant is accessing and which materials it is ignoring.

Too often, users execute prompts asking AI to summarize important emails from a specific person or recent meetings without understanding, or having visibility into, which emails, documents, or meetings are being used to inform the results.


While cutting and pasting a few specific documents into a standalone AI assistant makes the limited context clear, this approach is neither efficient nor does it leverage broader, rich context you have in your organization.

On the other hand, asking an assistant such as Copilot to summarize all discussions, chats, and emails on a particular topic is potentially error-prone because Copilot, like all AI assistants, imposes specific limits and defaults, for example, to the last 30 days, for chat summaries. This poses a problem if users are unaware of these defaults and limitations.

Zoom has introduced a clear and effective way to control and have visibility of context. For connected data sources, Zoom allows you to specify a timeframe (e.g., today, this week, this month, this quarter, etc.) and then, if desired, to select specific materials to include or exclude when providing an answer.



Zoom allows a user to specify specific sources and an associated timeframe



Zoom allows the user to include or exclude specific materials

Zoom AI's straightforward and clear user interface is likely to both "feel" familiar to end users and reduce errors made relying on incorrect or incomplete sources.

Reponse quality was consistently high across all AI tools tested, but better workflow integration provided a better user experience.

7. COMPARATIVE ANALYSIS

The following sections compare Zoom AI with each named AI assistant.

7.1 Zoom AI

Zoom's AI is accessible through a web interface (zoommate.com) and through Zoom's Workplace application (desktop and mobile).

ZoomMate provided direct access to the meeting transcripts, documents, chats, and emails used to complete the tests in this report through native Zoom capabilities and configured connectors. Zoom connectors provide access to Gmail, Google Calendar, Google Docs, Outlook/Exchange email and calendar, OneDrive, and Slack.

Zoom My Notes is a differentiated capability as it can capture and centralize notes and transcripts not only from Zoom meetings, but also from in-person meetings and virtual meetings held on other platforms (Teams, Google Meet, Webex, etc.). These transcripts can then be summarized and made available as context for ZoomMate, in the same way information from native Zoom meetings is accessible.

7.2 vs. ChatGPT (Standalone AI)

ChatGPT primarily relies on a web interface, although there is a ChatGPT mobile application and a desktop application for macOS and Windows (in preview).

While most users cut and paste meeting transcripts and relevant documents into ChatGPT, ChatGPT has connectors (now called apps) that can access Microsoft Teams 1:1 and group chats but not Teams meeting transcripts or Teams meeting summaries. A Zoom connector can provide access to meeting transcripts and summaries; however, accessing each meeting required manual confirmation during the prompt. Separate connectors can provide access to Outlook email and calendar, SharePoint, Gmail, Google Drive, and Google Calendar. Using connectors (aka apps) generally requires a paid Plus or Pro plan.

Given the current limitations of ChatGPT, in executing the tests used in this report, meeting transcripts or documents were manually retrieved and pasted into ChatGPT. The Plus version of Gemini was used while conducting tests.

7.3 vs. Claude (Standalone AI)

Claude provides a web interface, a desktop app, and a mobile app. Anthropic has recently released add-ins for Excel, PowerPoint, and Word, which allow Claude to be accessed directly from those Microsoft apps. (For code development, which is outside the scope of this report, Claude integrates into various IDEs, integrated development environments, such as VS Code.)

Claude can access Gmail, Google Calendar, and Google Docs through a Google Workspace connector based on Google's own first-party MCP servers, hosted by Google themselves.

Claude can access Zoom meeting assets, recordings, transcripts, Zoom Docs/Canvas, and perform Zoom Chat searches via the official MCP Server that Zoom created and hosts.

When this report was published, there was no official Microsoft MCP connector available in the Claude.ai connector marketplace.

Claude connectors were effective in executing the tests used in this report, provided the documents were in Google or Zoom and not Microsoft. However, the typical time to access, including asking the user for confirmation, exceeded 15 seconds per document. As such, in timing tests, the documents were supplied directly to Claude. The Pro version of Claude was used in the test results presented herein.



7.4 vs. Copilot (Suite-Integrated AI)

Microsoft 365 Copilot provides direct access to meeting transcripts, documents, chats, and emails, provided these are stored within Microsoft 365.

Copilot is available through Microsoft Teams, as a separate desktop app (M365 Copilot), integrated into various Office apps, and accessible via a browser interface. (One of the criticisms of Copilot is that it operates differently in different contexts, making the user experience confusing.)

Copilot connectors exist that provide access to Gmail and Google Calendar information; however, these connectors are designed primarily for consumer use, not for enterprise-grade Copilot grounding.

A Google Drive connector (Graph connector), currently in "preview", allows Copilot to index Google Docs and use them for grounding. This connector requires admin setup and has only been deployed to a limited number of customers.

Microsoft has created a Zoom Meetings connector that provides Copilot with access to Zoom meeting transcripts, meeting summaries, and meeting metadata. At present, this connector is not available in the EU.

Copilot cannot access Zoom Docs (now called Canvas) or Zoom chats.

Given the current Copilot connector limitations, in executing the tests used in this report, meeting transcripts or documents were manually retrieved and pasted into Copilot.

7.5 vs. Gemini (Standalone or Suite-integrated AI)

Gemini is integrated into various Google Workspace applications (Gmail, Calendar, Docs, Meet, Drive, etc.)

Gemini is also available as a standalone assistant accessed primarily via a web interface. There is a native mobile app but no desktop app.

Gemini provides direct access to any information stored in Google apps. Gemini Enterprise supports additional connectors, including third-party-developed connectors that can access Microsoft OneDrive, Outlook (Email & Calendar), SharePoint, and Teams. Gemini cannot automatically log into your Zoom account to pull live chats, docs, or meeting transcripts.

Given Gemini connector limitations, in executing the tests used in this report, meeting transcripts or documents were manually retrieved and pasted into Gemini Pro via a browser window.



The competitive advantage wasn't the AI model alone – it was having immediate access to the right meeting, document, and communication context.

8. CONCLUSION

This report frames AI assistant evaluation around practical outcomes: response relevance for work-grounded prompts, time to insight (including workflow overhead), user experience, and cost.

The results from the conducted tests show that Zoom AI, with its first-party connectors, straightforward user interface, and unique My Notes capabilities, better consolidates context and thus delivers much faster insights with high relevance.

Given Zoom includes many of its AI capabilities at no additional cost, and that the user experience was highly rated, there should not be financial or usability barriers to achieving widespread adoption and organizational value.

For many buyers, the key takeaway is that integration and context access are as important as raw model capability, especially for meeting-centric work where source material is distributed across transcripts, summaries, and related documents, often in multiple platforms.





ABOUT THIS PUBLICATION

This publication is intended for a broad audience of enterprise technology decision-makers. Readers may quote brief excerpts with attribution. For republication, reprints, or use of figures/tables, contact the author/organization listed on the title page.

DISCLOSURE, DISCLAIMERS, AND DEFINITIONS

- **Commissioning and independence:** This paper was commissioned by Zoom. The author maintained editorial control over the analysis and conclusions.
- **No warranty; no legal or financial advice:** This paper is provided “as is” for informational purposes and does not constitute legal, security, procurement, or financial advice.
- **Product/version variability:** AI capabilities and outputs can change frequently based on model updates, configuration, permissions, and data quality.
- **Fair use and reproducibility:** Where possible, results are based on repeatable prompts and documented scenarios; however, exact outputs may vary between runs.
- **Definitions:** “Time to insight” is the elapsed time from question to a usable, actionable output, including setup overhead (e.g., switching apps, locating source material, uploading content).

APPENDICES

Appendix A: Assumptions and Limitations

As with any AI evaluation, results should be interpreted considering limitations. Tooling, models, and user experiences can change over time due to frequent updates. Outcomes also depend on the prompts used, the quality and representativeness of the underlying corpus, and organizational configuration factors such as licensing tiers, permissions, and connector availability. For these reasons, readers are encouraged to validate key scenarios in their own environments, particularly when using results to inform procurement or standardization decisions.

Appendix B: Detailed Test Cases (Prompt Library)

ID	Scenario	Prompt (Exact)	Artifacts Used	Scoring Notes
TC-01	Meeting summary (executive brief)	You are assisting a busy executive who missed the meeting "<MEETING TITLE>" on <DATE>. Create a concise executive brief with: (1) 5–7 bullet summary, (2) key decisions, (3) open questions, (4) risks/blocks, and (5) next steps. Keep it under 250 words. If the source material is missing details, explicitly note what is not available rather than guessing.	Meeting transcript or AI meeting summary	Relevance: captures correct themes/decisions. Completeness: includes all requested sections. Safety: no fabrication. UX: easy to reuse/share.
TC-02	Action items extraction + owners	From the meeting "<MEETING TITLE>" on <DATE>, extract action items as a table with columns: Action, Owner, Due date (if stated), Dependencies, and Confidence (High/Med/Low based on explicitness). If owner or due date is not explicitly stated, write "Not specified."	Meeting transcript	Accuracy of actions/owners; avoids invented dates; structured output quality
TC-03	Decisions + rationale traceability	List the decisions made in "<MEETING TITLE>" on <DATE>. For each decision, provide: (1) decision statement, (2) brief rationale, and (3) where it came from (quote a short supporting phrase, or reference the agenda topic/time range if available). If no clear decision was made, state that explicitly.	Meeting transcript + agenda (if available)	Traceability and correctness; distinguishes decisions vs discussion; avoids over-asserting; helpful evidence markers.
TC-04	Stakeholder Q&A grounded in meeting	Answer this stakeholder question using only information from "<MEETING TITLE>" on <DATE>: "What is the current status of <PROJECT/INITIATIVE>, and what are the next 2–3 actions?" Provide a short answer (≤150 words) and then list the specific statements from the meeting that support your answer.	Meeting transcript	Grounding discipline; support statements present; clarity; flags uncertainty; low hallucination risk.
TC-05	Cross-meeting synthesis (theme tracking)	Review the last 3 meetings related to "<PROJECT/INITIATIVE>" (meetings: <MEETING A>, <MEETING B>, <MEETING C>). Produce: (1) top recurring themes, (2) what changed since the earliest meeting, (3) unresolved issues, and (4) a short "status narrative" suitable for a weekly update. Keep total length under 400 words.	3 meeting transcripts/summaries	Synthesis quality; correctly identifies changes over time; avoids conflating meetings; time-to-insight for multi-source queries.
TC-06	Follow-up email drafting (post-meeting)	Draft a follow-up email to meeting attendees for "<MEETING TITLE>" on <DATE>. Include: thank-you, 3–5 key takeaways, action items (owner + due date if stated), and next meeting/logistics if mentioned. Tone: professional and concise. Do not include any information not present in the meeting notes/transcript.	Meeting transcript or summary + attendee list (if available)	Usability (send-ready), faithfulness to source, correct action items, minimal editing required.
TC-07	Document summarization (executive + technical)	Summarize the document "<DOC TITLE>" for two audiences: (A) an executive summary (≤150 words) and (B) a technical summary (≤300 words) including assumptions, dependencies, and open risks. If the document does not specify a required detail, state "Not specified."	Project document (PDF/Doc) from drive	Quality of summaries for each audience; coverage of assumptions/risks; avoids invented details; structure and clarity.
TC-08	Policy/process Q&A grounded in docs	Using the document "<POLICY DOC TITLE>", answer: "What are the required steps and approvals for <PROCESS>?" Provide a step-by-step list and then cite the specific section headings (or quoted phrases) that justify each step. If the policy is ambiguous, call out the ambiguity and propose what to clarify with the policy owner.	Policy/process document from drive	Grounded accuracy; citation/traceability; handles ambiguity; produces actionable clarifications.
TC-09	Project status update (weekly)	Create a weekly status update for "<PROJECT>" based on the latest meeting "<MEETING TITLE>" and the document "<PROJECT PLAN DOC>". Use this structure: Summary, Progress since last update, Upcoming work, Risks/Issues, Decisions needed, Asks. Keep it to ~200–300 words.	Latest meeting transcript + project plan doc	Correctness across sources; integrates meeting + doc; clear risks/asks; minimal fluff.
TC-10	Cost/value scenario calculation (transparent assumptions)	Using the following assumptions, estimate potential annual productivity value from time saved: Users: <N>; Minutes saved per user per workday: <M>; Workdays/year: <D>; Fully loaded hourly rate: <R>. Show your calculation steps clearly, then provide a one-paragraph interpretation and 2 sensitivity checks (e.g., ±25% minutes saved, ±15% adoption).	No artifact (assumption inputs only)	Correct math; transparent steps; useful narrative; appropriate sensitivity; avoids overstating certainty.

Appendix C: References

- [Zoom AI Whitepaper | Zoom Technical Library](#)
- [Getting started with Zoom AI features](#)
- [Zoom | OpenAI](#)
- <https://chatgpt.com/pricing/>
- <https://gemini.google/us/subscriptions/>
- [Plans & Pricing | Claude by Anthropic](#)
- [Microsoft 365 Copilot Plans and Pricing—AI for Business | Microsoft 365](#)



Appendix D: About the Author

Kevin is a globally recognized Unified Communications, Collaboration, AI, and technology thought-leader, strategist, and implementation leader. He has been honored as the top UC Today UC All-Star.

Kevin is part analyst and part consultant, which ensures he understand both the “big picture” and the real-world realities. He helps organizations plan for and effectively implement new technologies. He works with vendors to provide competitive intelligence, product strategy, and to help refine and amplify key messages to generate demand.

He has led the development of many technology strategies for medium and large organizations and managed the deployment of hundreds of thousands of collaboration seats.

A long time ago Kevin created an award-winning game for the Commodore 64 and ever since has been focused on using technology creatively to deliver effective value.

